CN5000 Maintenance and Troubleshooting

Guide

Rev. 1.7
April 2026

Table of Contents

Revision History

Date	Rev	Description
Apr 2026	1.7	• Removed "Monitoring Fabric Performance" using the opatop TUI. TUI no longer supported. • Updated Section 3.2.5.2 "Upgrading the OPX Software", "Create the Repository File" and "Import the GPG Key" steps.
Mar 2026	1.6	• Updated "Upgrading the OPX Software", Step 6 for resolving Ubuntu dependency conflicts.
Jan 2026	1.5	• Added new section, "Remove a Liquid-Cooled Switch".
Nov 2025	1.4	• Added note in "Update the Switch Firmware" that BMC firmware may require multiple (2) updates/reboots. • Updated "Update the OPX Software" to include new Ubuntu instructions. • Updated "Remove OPX Software from a Host" to include new Ubuntu instructions.
Oct 2025	1.3	• Updated "LinkPolicy" to include important note on forcing a link policy change.
Sep 2025	1.2	• Updated "Update the Switch Firmware" to simplify steps and include both Push and Pull Methods. • Moved hardware LED information to *CN5000 Product Family Description Guide*. Reference provided in "Identifying Issues from LEDs". • Added new section, "Ports Speed" to ensure cable-supported speeds are configured for links.
Aug 2025	1.1	• Added "Monitoring SuperNIC and AOC Temperatures" for in-band temperature sensing. • Added magnified view of Switch LEDs to "Switch LEDs". • Added "Failed to Start Fan Sensor in Liquid-Cooled Switches" to notify users of a fan sensor error message on liquid-cooled Switches that can be ignored.
Aug 2025	1.0	• Initial release.

Preface

Cornelis Networks delivers the world's highest performance scale-out networking solutions for AI and HPC data centers. Our differentiated architecture seamlessly integrates hardware, software, and system level technologies to maximize the efficiency of GPU, CPU, and accelerator-based compute clusters at any scale. Our solutions drive breakthroughs in AI and HPC workloads, empowering our customers to push the boundaries of innovation.

This guide is part of the CN5000 documentation set. Omni-Path is an end-to-end solution consisting of CN5000 SuperNICs, CN5000 Switches, Director Class Switches, and fabric management software and tools.

License Agreements

Cornelis software and firmware are provided under one or more license agreements. Refer to the license agreement(s) provided with the software for specific details. Do not install or use the software until you have carefully read and agreed to the terms and conditions of the license agreement(s). By loading or using the software, you agree to the terms of the license agreement(s). If you do not wish to so agree, do not install or use the software.

Technical Support

Technical support for Cornelis products is available 24 hours a day, 365 days a year:

- Website: Cornelis Technical Support
- Email: support@cornelisnetworks.com
- Phone: +1 484-497-9665

Best Practices

- To obtain the most recent functional and security updates, Cornelis recommends that you download and install the latest versions of the CN5000 Omni-Path software and firmware. Software and Documentation are found in the Cornelis Customer Center.
- To ensure optimal performance in your CN5000 fabric, it is important to review the Cornelis performance tuning recommendations in the *CN5000 Performance Tuning Guide*.

Document Conventions

Omni-Path documentation uses the following standard conventions:

- **Note:** Provides additional information.
- **Caution:** Indicates the presence of a hazard that has the potential of causing damage to data or equipment.
- **Warning:** Indicates the presence of a hazard that has the potential of causing personal injury.
- Text in blue indicates a hyperlink to a figure, table, or section in this guide. Links to websites are also shown in blue. For example:
  - See License Agreements for more information.
  - For more information, visit Cornelis Networks.
- Text in **bold** indicates user interface elements such as menu items, buttons, check boxes, key names, keystrokes, or column headings. For example:
  - Click the **Start** button, point to **Programs**, point to **Accessories**, and then click **Command Prompt**.
  - Press **CTRL+P** and then press the **UP ARROW** key.
- Text in `Courier` font indicates a file name, directory path, or command line text. For example:
  - Enter the following command: `port show`
- Preformatted text in `Courier` font with a gray background indicates a block of code.

```
Welcome to Cornelis Networks shell
Use 'tab' for autocomplete and up/down arrow to see command history.
```

- Text in *italics* indicates terms, emphasis, variables, or document titles. For example:
  - Refer to *CN5000 Fabric Installation Guide* for details.
- Both the CN5000 Omni-Path Fabric and standard InfiniBand (IB) can send IP traffic over the fabric, or *IPoFabric*. In this document, it may also be referred to as *IP over IB* or *IPoIB*. From a software point of view, IPoFabric behaves the same way as IPoIB.
- InfiniBand interfaces follow different naming conventions depending on the operating system distribution. Throughout this document set, the interfaces will be referenced as "ib0" and "ib1" for consistency. Your system may use different naming schemes based on its specific OS implementation.
- Many of the acronyms in this document link to the glossary.

NOTE

In this document, references to "HFI" refer to the CN5000 SuperNIC.

CN5000 Omni-Path Documentation Set Overview

The following documents comprise the CN5000 Omni-Path documentation set (in order of recommended reading).

- *CN5000 Product Family Description Guide*

  This document describes the CN5000 Omni-Path products and their features at a high level.

- *CN5000 Fabric Installation Guide*

  This guide contains instructions for the hardware installation, software installation, initial configuration, and verification tasks required for the successful installation of a CN5000 Fabric.

- *CN5000 Topologies and Routing Guide*

  This document describes the different topologies and routing methods that can be used with the CN5000 Omni-Path Fabric.

- *CN5000 Provisioning and Security Guide*

  After you have installed and configured your CN5000 hardware and software and set up your topology, this document provides you with the information to provision your virtual fabrics and set up security, if required.

- *CN5000 Performance Tuning Guide*

  This document describes the BIOS settings and parameters that have been shown to optimize performance on Omni-Path. If you are interested in benchmarking the performance of your system, these tips may help you obtain better performance.

- *CN5000 Maintenance and Troubleshooting Guide*

  This document provides you with the information to maintain and troubleshoot your CN5000 Omni-Path products and fabric.

- *CN5000 Commands Guide*

  This reference document describes the CLI environment and commands used for managing the Omni-Path hardware and software. This document is a companion to all other task-oriented documentation. Refer to it when you need details pertaining to CLI commands and configuration files.

1. Introduction

The CN5000 Omni-Path Fabric is the ideal HPC fabric for interconnecting resources, using a scalable, 400 Gbps, low-latency fabric, delivering an exceptional set of high-speed networking features and functions.

This document provides you with the information to maintain and troubleshoot your CN5000 Omni-Path products and fabric.

1.1. Intended Audience

This document is intended for hardware installers, fabric administrators, and other personnel with similar qualifications.

1.2. Layout and Document Organization

This document provides detailed information and procedures for maintaining and troubleshooting your Omni-Path system.

This document is organized as follows:

- Section 2 "Monitoring" provides the stand up and day-to-day tasks and information to help you ensure your CN5000 Omni-Path system is working properly.
- Section 3 "Maintenance" provides common procedures and information to maintain your CN5000 Omni-Path system.
- Section 4 "Troubleshooting" provides basic troubleshooting methods and information for your CN5000 Omni-Path system.
- Section 5 "Working with Cornelis Technical Support" provides expert assistance for a variety of services including technical issues and questions, warranty, and returns.

2. Monitoring

This chapter provides the stand up and day-to-day tasks and information to help you ensure your CN5000 Omni-Path system is working properly.

NOTE

In this document, references to "HFI" refer to the CN5000 SuperNIC.

2.1. Setting Up Remote Logging

To avoid losing log information in the event of a hardware failure, Cornelis recommends that you configure a remote syslog server.

If you have not already set up remote logging, refer to the *CN5000 Fabric Installation Guide* for details.

2.2. Monitoring Logs and Events

Logs and events provide valuable insights into system behavior, security events, and performance. They can enable faster troubleshooting and problem resolution.

2.2.1. Fabric Manager Monitoring

The Fabric Manager logs events to assist with fabric debugging.

The following lists some of the common Fabric Manager message log information that you should look for:

- Normal Sweeps
- Components that are constant
- No Link Quality Indicator (LQI) errors

The following steps provide high-level instructions for monitoring the fabric/system using the Fabric Manager.

1. Ensure `/var/log/message` is clean.

   NOTE

   If opafm logs have been redirected elsewhere, check that directory.

   A healthy fabric should have sweeps every five minutes unless there is an event in the fabric, such as node reboot or adding or removing connections.

   You may see some retries in the “retries” counter, but the number should be in single or low-double digits. The sweep time should be 0.2 – 0.5 seconds (depending on the size of the fabric—in large setups, it may be greater than 1 second).

Example output:

```
Feb 11 08:45:57 opahsx151 fm0_sm[46537]: PROGR[topology]: SM: topology_main: TT:
DISCOVERY CYCLE START - REASON: Scheduled sweep interval
Feb 11 08:45:57 opahsx151 fm0_sm[46537]: PROGR[topology]: SM: sm_set_local_port_pkey:
sm pkey table already set
Feb 11 08:45:57 opahsx151 fm0_sm[46537]: PROGR[topology]: SM: topology_main:
DISCOVERY CYCLE END. 2 SWs, 30 HFIs, 30 end ports, 68 total ports, 1 SM(s), 148
packets, 0 retries, 0.238 sec sweep

<after 5 minutes>

Feb 11 08:50:17 opahsx151 fm0_sm[46537]: PROGR[topology]: SM: topology_main: TT:
DISCOVERY CYCLE START - REASON: Multicast group Membership change.
Feb 11 08:50:17 opahsx151 fm0_sm[46537]: PROGR[topology]: SM: sm_set_local_port_pkey:
sm pkey table
Feb 11 08:50:17 opahsx151 fm0_sm[46537]: PROGR[topology]: SM: topology_main:
DISCOVERY CYCLE END. 2 SWs, 30 HFIs, 30 end ports, 68 total ports, 1 SM(s), 149
packets, 0 retries, 0.217 sec sweep
```

Any messages with **WARN** or **ERROR** should be investigated.

2. Show only the most recent Fabric Manager log entries, and continuously print new entries so they are appended to the journal.

```
journalctl -f -u opafm
```

2.2.2. Monitoring SuperNIC and AOC Temperatures

The SuperNIC allows for reporting of its current temperature as well as the temperature of any active optical cables present. This can be done using the `tempsense` file:

```
[root@Genoa165 ~]# cat /sys/class/infiniband/hfi1_0/tempsense
ASIC 55.250
QSFP1 none
QSFP2 52.625
```

The temperatures listed here are in degrees Celsius.

2.3. Common Fabric Monitoring Tools

This section describes the common fabric monitoring tools and commands to check cluster and software status.

2.3.1. opafabricinfo

Use `opafabricinfo` to monitor fabric components, using the first active port on the given local host to perform its analysis.

From this command, you can determine information such as which host is running the primary FM and how many inter-switch links (ISLs) there are.

```
# opafabricinfo
Fabric 0:0 Information:
SM: opahsx151 hfi1_0 Guid: 0x0011750101671d2c
State: Master
Number of HFIs: 30
Number of Switches: 2
Number of Links: 33
Number of HFI Links: 30             (Internal: 0   External: 30)
Number of ISLs: 3                   (Internal: 0   External: 3)
Number of Degraded Links: 0         (HFI Links: 0   ISLs: 0)
Number of Omitted Links: 0          (HFI Links: 0   ISLs: 0)
--------------------------------------------------------------------------------
```

For additional details, refer to *CN5000 Commands Guide*, "opafabricinfo".

2.3.2. opareport

Use `opareport` to generate reports about the current state of the fabric and output snapshot files for later use in debugging.

Common issues to analyze are:

- Cable health, including link quality
- Nodes that are Inactive, Priority, or Elevated Priority when controlling failover for SM, PM

The following provides examples for looking into cable health, through errors and slow links.

NOTE

Before running `opareport` to look for errors (`-o`), you should clear all the counters using `opareport -o none --clearall` or `opareport -o none -Ca`.

```
opahsx41:~ # opareport -o none --clearall
Getting All Node Records...
Done Getting All Node Records
Done Getting All Link Records
Done Getting All Cable Info Records
Done Getting All SM Info Records
Done Getting vFabric Records
Clearing Port Counters

Configured Counters to Clear:
XmitData
RcvData
XmitPkts
```

```
RcvPkts
MulticastXmitPkts
MulticastRcvPkts
UncorrectableErrors
LinkDowned
RcvErrors
ExcessiveBufferOverruns
FMConfigErrors
LinkErrorRecovery
LocalLinkIntegrityErrors
RcvRemotePhysicalErrors
XmitConstraintErrors
RcvConstraintErrors
RcvSwitchRelayErrors
XmitDiscards
CongDiscards
RcvFECN
RcvBECN
MarkFECN
XmitTimeCong
XmitWait
XmitWastedBW
XmitWaitData
RcvBubble
Clearing Port Counters...
Done Clearing Port Counters
Cleared 68 Ports on 32 Nodes
```

- To check for **cable errors and slow links**, use `opareport -o errors -o slowlinks`.

  Cables with LinkQualityIndicator (LQI) "3" or less should be troubleshot or replaced. For more information on LQI, refer to Section 4.2.5.2.1 “Link Quality Indicator (LQI)”.

```
opareport -o errors
Getting All Node Records...
Done Getting All Node Records
Done Getting All Link Records
Done Getting All Cable Info Records
Done Getting All SM Info Records
Done Getting vFabric Records
Getting All Port Counters...
Done Getting All Port Counters
Links with errors > threshold Summary

Configured Thresholds:
  LinkQualityIndicator           3
   UncorrectableErrors            1
   LinkDowned                     3
   RcvErrors                      1
   ExcessiveBufferOverruns        1
   FMConfigErrors                 1
   XmitConstraintErrors           10
   RcvConstraintErrors            10
   CongDiscards                   100
Rate NodeGUID          Port Type Name
400g 0x001175010170c572   1 FI   ime06 hfi1_0
```

```
   LinkQualityIndicator: 3 Below Threshold: 4
<->  0x00117501020c4bdb  10 SW  sw_ddn_r25u41
17688 of 17688 Links Checked, 1 Errors found
```

- To check for **slow links** (links on which lanes, either RX or TX, are running at speed less than "4"), use `opareport -o slowlinks`.

  The following example shows two ends of a cable with “3” specifying one degraded lane TX/RX.

```
opareport -o slowlinks
Getting All Node Records...
Done Getting All Node Records
Done Getting All Link Records
Done Getting All Cable Info Records
Done Getting All SM Info Records
Done Getting vFabric Records
Links running slower than expected Summary

Links running slower than expected:
Rate NodeGUID          Port Type Name
    Active                          Enabled
    Lanes, Used(Tx), Used(Rx), Rate,    Lanes,   DownTo,  Rates
-------------------------------------------------------------------------------
400g 0x001175010170b5a9   1 FI   ddn-mon02 hfi1_0
    4     3          4         100Gb     4        3,4      100Gb
<->  0x00117501020c4b3e  15 SW  sw_ddn_r25u42
    4     4          3         100Gb     1,2,3,4  3,4      100Gb
```

For additional details and the list of available report types, refer to *CN5000 Commands Guide*, “opareport”.

2.3.3. opainfo

Use `opainfo` to report on the status of the local SuperNICs.

The following provides an example output for a single port SuperNIC.

```
$ opainfo
hfi1_0:1                          PortGID:0xfe80000000000000:0011750101575fec
   PortState:     Active
   LinkSpeed      Act: 100Gb         En: 100Gb
   LinkWidth      Act: 4             En: 4
   LinkWidthDnGrd ActTx: 4  Rx: 4    En: 3,4
   LCRC           Act: 14-bit        En: 14-bit,16-bit,48-bit       Mgmt: True
   LID: 0x00000001-0x00000001        SM LID: 0x00000001 SL: 0
   QSFP: PassiveCu,   2m TE Connectivity   P/N 2821076-2         Rev B
   Xmit Data:                  5 MB Pkts:                28742
   Recv Data:                 17 MB Pkts:                28969
   Link Quality: 5 (Excellent)
```

For additional details, refer to *CN5000 Commands Guide*, “opainfo”.

2.3.4. opatop

Use the Performance Monitor Tool `opatop` to *drill* down from a high-level, fabric-wide view to an individual port view.

From this tool, you can determine when an issue occurred at the high-level and drill down to find the offending port.

The following provides an example Summary screen.

```
opatop: Img: 10s @ Wed Mar 12 16:38:13 2025, Live
Summary:  SW:     1 Ports: SW:     5  HFI:     4       Link:     4
          SM:     1 Node NRsp:     0 Skip:     0 Port NRsp:     0 Skip:     0
                    AvgMBps   MinMBps   MaxMBps   AvgKPps   MinKPps   MaxKPps
0 All         Int         0         0         0         0         0         0
      Integ:min  Congst:min  SmaCong:min  Bubble:min  Secure:min  Routing:min
1 HFIs        Snd         0         0         0         0         0         0
              Rcv         0         0         0         0         0         0
      Integ:min  Congst:min  SmaCong:min  Bubble:min  Secure:min  Routing:min
2 SWs         Int         0         0         0         0         0         0
              Snd         0         0         0         0         0         0
              Rcv         0         0         0         0         0         0
      Integ:min  Congst:min  SmaCong:min  Bubble:min  Secure:min  Routing:min

    Master-SM: LID: 0x0001 Port: 1   Priority: 0  State: Master
               Name: hds1fnb1051 hfi1_0
               PortGUID: 0x00117501010AA4D8
 Secondary-SM: none



Quit up Live/rRev/fFwd/time/bookmrked Bookmrk Unbookmrk ?help |
sS Pmcfg Imginfo View 0-n:
```

2.4. Topology Verification

The FastFabric Tools provide several ways to assist you in validating a running fabric's configuration and layout against a predefined/expected topology. This verification process can help you to identify issues including detecting missing cables, hosts, or switches; or verifying that cables are in the correct places. You can run topology verification tools during the initial fabric startup or at any other time after a fabric is configured. You can also set up the Fabric Manager to verify the topology every time a sweep of the fabric is done (using the Fabric Manager's predefined topology verification feature).

For more details, refer to the *CN5000 Fabric Installation Guide*, “Verify Topology”.

2.5. Performance Monitoring

This section provides information on the performance of the fabric using data from the Performance Manager parameters.

2.5.1. Performance Manager Parameters

The following tables describe `opafm.xml` parameters that can be used in the `Pm` subsection of either the `Common` or `Fm` sections.

NOTE

Any parameter that can be used in the `Common.Shared` section can also be used in the `Common.Pm` or `Fm.Pm` sections.

2.5.1.1. PM Controls

The parameters shown in the following table set up the options for when and how the PM monitors the fabric.

Table 1. PM Parameters

Parameter	Default Value	Description
ServiceLease	60	ServiceRecord lease with SA in seconds.
SweepInterval	10	The PM constantly sweeps and computes fabric statistics. If the `SweepInterval` is set to `0`, the PM will not perform sweeps. But instead if queried it will do an immediate PMA operation. Tools such as `opatop` require PM `SweepInterval` be non-zero. The default in the sample FM configuration file is ten seconds. However when upgrading from previous FM releases, if an FM configuration file is used without this value specified, a default of `0` is used. This permits upgrades to operate in a mode comparable to the existing configuration and requires specific user action to enable the new 6.0 and above PM features.
MaxClients	3	The maximum number of concurrent PA client applications (for example, `opareport`, `opatop`, `oparfm`) running against the same PM/PA.
TotalImages FreezeFrameImages	10 5	The PM can retain recent fabric topology and performance data. Each such dataset is referred to as an Image. Images allow for access to recent history and/or Freeze Frame by clients. Each image consumes memory, so care must be taken not to take an excessive amount of memory, especially for larger fabrics. `TotalImages` - total images for history and freeze `FreezeFrameImages` - max unique frozen images

Parameter	Default Value	Description
FreezeFrameLease	60	If a PA client application hangs or dies, after this set time, all its frozen images will be released. Specified in seconds

2.5.1.2. PA Category Parameters

2.5.1.2.1. PM Thresholds

PM Threshold parameters are set for each PM category. Exceeding the values for each category will result in a log warning. You can set a PM threshold value to 0 to ignore the given class of errors.

The parameters shown in the following table set the thresholds for each category and exceeding these values for each category will print a log warning. 0 causes the given class of errors to be ignored.

Table 2. Threshold Parameters

Parameter	Default Value	Description
Integrity	100	Threshold for logging a warning indicating a possible error condition.
Congestion	100	Threshold for logging a warning indicating a possible error condition.
SmaCongestion	100	Threshold for logging a warning indicating a possible error condition.
Bubble	100	Threshold for logging a warning indicating a possible error condition.
Security	10	Threshold for logging a warning indicating a possible error condition.
Routing	100	Threshold for logging a warning indicating a possible error condition.

- **TrapThreshold**

  TrapThreshold is configured to monitor the number of traps per minute for a port. If a given port exceeds the threshold value, the port is disabled as unstable.

  You can set the TrapThreshold to 0 to disable this feature.

- **TrapThresholdMinCount**

  TrapThresholdMinCount defines the minimum number of traps required to reach the TrapThreshold rate. For example, if TrapThreshold is set to 10 traps per minute and TrapThresholdMinCount is set to 5, the port is disabled after 5 traps are received at the rate of 10 traps per minute (that is, 5 traps in 30 seconds).

  The TrapThresholdMinCount parameter value must be greater than 2; the default value is 10.

If you set the TrapThreshold to 0, the TrapThresholdMinCount parameter is ignored.

Larger values of TrapThresholdMinCount increase the accuracy of detecting the trap rate, but also increase the time between a trap surge and the SM disabling a port.

Very small values of TrapThresholdMinCount can result in a port being disabled after a few traps.

2.5.1.2.2. Threshold Exceeded Message Limit

PM Threshold Exceeded Message Limit parameters limit how many ports that exceed their PM thresholds are logged per sweep. These parameters can help avoid excessive log messages when extreme fabric problems occur.

The parameters shown in the following table limit how many ports which exceed their PM Thresholds are logged per sweep. These can avoid excessive log messages when extreme fabric problems occur. These parameters can be used in the `ThresholdsExceededMsgLimit` section.

Table 3. PM ThresholdsExceededMsgLimit Parameters

Parameter	Default Value	Description
Integrity	10	Maximum ports per PM Sweep to log if exceeds configured threshold. A value of `0` suppresses logging of this class of threshold exceeded errors.
Congestion	0	Maximum ports per PM Sweep to log if exceeds configured threshold. A value of `0` suppresses logging of this class of threshold exceeded errors.
SmaCongestion	0	Maximum ports per PM Sweep to log if exceeds configured threshold. A value of `0` suppresses logging of this class of threshold exceeded errors.
Bubble	0	Maximum ports per PM Sweep to log if exceeds configured threshold. A value of `0` suppresses logging of this class of threshold exceeded errors.
Security	10	Maximum ports per PM Sweep to log if exceeds configured threshold. A value of `0` suppresses logging of this class of threshold exceeded errors.
Routing	10	Maximum ports per PM Sweep to log if exceeds configured threshold. A value of `0` suppresses logging of this class of threshold exceeded errors.

2.5.1.2.3. Integrity Weights

PM Integrity Weight parameters control the weights for the individual counters. These are combined to form the Integrity count. You can set a PM integrity weight parameter value to 0 to ignore the counter.

The parameters shown in the following table control the weights for the individual counters which are combined to form the `Integrity` count. These parameters can be used in the `IntegrityWeights` section.

Table 4. PM IntegrityWeights Parameters

Parameter	Default Value	Description
LocalLinkIntegrityErrors	0	Weight for LocalLinkIntegrityErrors counter. 0 causes the counter to be ignored.
RcvErrors	100	Weight for RcvErrors counter. 0 causes the counter to be ignored.
ExcessiveBufferOverruns	100	Weight for ExcessiveBufferOverruns counter. 0 causes the counter to be ignored.
LinkErrorRecovery	0	Weight for LinkErrorRecovery counter. 0 causes the counter to be ignored.
LinkDowned	25	Weight for LinkDowned counter. 0 causes the counter to be ignored.
UncorrectableErrors	100	Weight for UncorrectableErrors counter. 0 causes the counter to be ignored.
FMConfigErrors	100	Weight for FMConfigErrors counter. 0 causes the counter to be ignored.
LinkQualityIndicator	40	Weight applied to the LQI normalization equation. Calculation is 2^(5-LQI)-1. 0 causes the counter to be ignored.
LinkWidthDowngrade	100	Weight applied to the LWD normalization equation. Calculation is equal to the number of active lanes down: LinkWidth.Active - LinkWidthDowngrade.RxActive 0 causes the counter to be ignored.

2.5.1.2.4. Congestion Weights

The parameters shown in the following table control the weight to use for each individual counter when computing congestion, which are combined to form the `Congestion` count. Integrity errors can also cause congestion.

Percentage (Pct) means the specified counter is divided by the appropriate data transfer counter, then normalized to a predetermined range before being weighted and summed into the category.

Table 5. PM Congestion Weights Parameters

Parameter	Default Value	Description
XmitWaitPct	0	Weights for XmitWeight counter divided by an associated data transfer counter and normalized to a predetermined range of 0.01% to 1%. 0 causes the counter to be ignored.

Parameter	Default Value	Description
CongDiscards	100	Weights for SwPortCongestion counter. 0 causes the counter to be ignored.
RcvFECNPct	5	Weights for RcvFECN counter divided by an associated data transfer counter and normalized to a predetermined range of 0.1% to 10%. 0 causes the counter to be ignored.
RcvBECNPct	1	Weights for RcvBECN counter divided by an associated data transfer counter and normalized to a predetermined range of 0.1% to 10%. 0 causes the counter to be ignored.
XmitTimeCongPct	25	Weights for XmitTimeCongestion counter divided by an associated data transfer counter and normalized to a predetermined range of 0.1% to 10%. 0 causes the counter to be ignored.
MarkFECNPct	25	Weights for MArkFECN counter divided by an associated data transfer counter and normalized to a predetermined range of 0.1% to 10%. 0 causes the counter to be ignored.

2.5.1.3. PM Sweep Operation Control

The parameters shown in the following table control the operation of the PM during each sweep.

Table 6. PM Sweep Parameters

Parameter	Default Value	Description
Resolution		Resolution determines the number of LocalLinkIntegrity or LinkErrorRecovery errors that must occur before the PMA will include them in the ErrorCounterSummary. Most counters are 64 bits wide and are expected to not ever saturate in the uptime of a port.
.LocalLinkIntegrity	8000000	
.LinkErrorRecovery	100000	
ErrorClear	7	This controls when the PM clears PMA Error counters 0 = clear when non-zero 1 = clear when 1/8 of individual counters max 2 = clear when 2/8 of individual counters max ... 7 = clear when 7/8 of individual counters max
ClearDataXfer	0	Enable clearing of Data Transfer Counters.
Clear64bit	0	Enable clearing of 64-bit Error Counters.
Clear32bit	1	Enable clearing of 32-bit Error Counters.
Clear8bit	1	Enable clearing of 8-bit Error Counters.

Parameter	Default Value	Description
ProcessHFICounters	1	Enable processing (sweeping) of SuperNIC Counters.
ProcessVLCounters	1	Enable processing (sweeping) of VL Counters.
PmaBatchSize	2	Maximum concurrent PMA requests the PM can have in flight while querying the PMAs in the fabric.
MaxParallelNodes	10	Maximum nodes to concurrently issue parallel requests to a given PMA.
MaxAttempts RespTimeout MinRespTimeout	3 250 35	The PM will spend up to `RespTimeout * MaxAttempts` per packet. These allow two modes of operation. When `MinRespTimeout` is non-zero, the PM will start with `MinRespTimeout` as the time-out value for requests and use multiples of this value for subsequent attempts if there is a time-out in the previous attempt. PM will keep retrying until the cumulative sum of time-outs for retries is less than `RespTimeout` multiplied by `MaxAttempts`. This approach is recommended and will react quickly to lost packets while still allowing adequate time for slower PMAs to respond. When `MinRespTimeout` is zero, upon a time-out, up to `MaxAttempts` are attempted with each attempt having a time-out of `RespTimeout`. This approach is provided for backward compatibility with previous PM versions.
SweepErrorsLogThreshold	10	Maximum number of PMA node or Port warning messages to output per sweep with regard to nodes that cannot be properly queried.

2.5.1.4. PM Overrides of the Common.Shared Parameters

The `Common.Shared` parameters can be overridden in the PM using the parameters described in the following table.

Table 7. Additional PM Parameters

Parameter	Default Value	Description
LogLevel	2	**NOTE:** Overrides the `Common.Shared` LogLevel Settings. Sets log level option for PM: • 0 = disable the vast majority of logging output • 1 = fatal, error, warn (syslog CRIT, ERR, WARN) • 2 = +notice, INFIINFO (progress messages) (syslog NOTICE, INFO) • 3 = +INFO (syslog DEBUG) • 4 = +VERBOSE and some packet data (syslog DEBUG) • 5 = +debug trace info (syslog DEBUG) This parameter is ignored for the Embedded Fabric Manager. For information on configuring chassis logging options, refer to the *CN5000 Commands Guide*, "log".
LogFile		**NOTE:** Overrides the `Common.Shared` LogLevel Settings. Sets log output location for PM. By default (or if this parameter is empty) log output is accomplished using syslog. However, if a LogFile is specified, logging is done to the given file. LogMode further controls logging. This parameter is ignored for the Embedded Fabric Manager. For information on configuring chassis logging options, refer to the *CN5000 Commands Guide*, "log".
SyslogFacility	Local6	**NOTE:** Overrides the `Common.Shared` LogLevel Settings. For the Host Fabric Manager, controls what syslog facility code is used for log messages. Allowed values are: auth, authpriv, cron, daemon, ftp, kern, local0-local7, lpr, mail, news, syslog, user, or uucp. For the Embedded Fabric Manager, this parameter is ignored.
ConfigConsistencyCheckLevel	2	Controls the Configuration Consistency. Check for PM. If specified for an individual instance of PM, will override Shared settings. Checking can be completely disabled or can be set to take action by deactivating Secondary PM if configuration does not pass the consistency check criteria. • 0 = disable Configuration Consistency Checking • 1 = enable Configuration Consistency Checking without taking action (only log a message) • 2 = enable Configuration Consistency Checking and take action (log message and shutdown Secondary PM)
Priority	0	0 to 15, higher wins.
ElevatedPriority		0 to 15, higher wins.

2.5.1.5. PM Short-Term History PM Parameters

The following parameters are used to enable and customize the off-loading of RAM-resident images to disk in order to preserve counter history for a long period of time.

Table 8. Short-Term History PM Parameters

Parameter	Default Value	Description
Enable	1	Enable Short-Term History.
StorageLocation	`/var/usr/lib/opa-fm/pm0_pahistory`	The absolute path where the history files will be stored. StorageLocation is a PM instance-specific parameter (similar to TcpPort for the FE). The default value is `/var/usr/lib/opa-fm/pm0_pahistory` (where '0' is the number of the instance). If there are multiple instances of the PM, then `StorageLocation` must be unique for each of them.
TotalHistory	24	The total number of hours of history that will be stored. This is a limit on the total amount of data stored, and not a limit on a file's age. PM downtime does not count toward the total time.
ImagesPerComposite	3	Determines how many images will be compounded into a single image as part of writing to the file. A higher number will save disk space but will result in lower data granularity. Must not be 0.
MaxDiskSpace	1024	A cap on how much disk space (in MiB) the short-term history is allowed to use. If this size is exceeded, the oldest files will be deleted to save space.
CompressionDivisons	8	Determines how many divisions will be used to concurrently compress or decompress data. Recommend less than or equal to number of processing cores of the management node, which must not exceed 32.

2.5.1.6. PM/PA Fail-over Parameters

The following parameters are used to enable and control the PM Fail-over extension.

Table 9. PM/PA Fail-over Parameters

Parameter	Default Value	Description
ImageUpdateInterval	5	`ImageUpdateInterval` is defined as the interval at which the Primary PM updates Standby PMs with an image (when an image is available); otherwise, image updates occur as they are available. Note that if multiple Standby PMs exist, the Primary PM updates each Standby concurrently during the `ImageUpdateInterval`. In the event PA Short-Term History is enabled on all PMs, the Primary PM will update all Standby PMs with disk-resident images once all RAM-resident images have been updated. `ImageUpdateInterval` must be less than the PM `SweepInterval` in order for the Primary PM to be able to send images fast enough to keep up with new images as well as catch up on older images. If `ImageUpdateInterval` is greater than `SweepInterval`, then `ImageUpdateInterval` is set equal to `SweepInterval` and a warning message is logged. Setting `ImageUpdateInterval` to 0 turns off the transfer of images to Standby PMs. The default value for `ImageUpdateInterval` is based upon the (`SweepInterval` / 2) rounded down to the nearest integer; must be at least 1.

2.5.1.7. Additional PM Parameters for Debug and Development

The PM supports the parameters in the following table to aid diagnosis and debugging. Only use these parameters under the direction of your support representative.

Table 10. PM Debug Parameters

Parameter	Default Value	Description
Debug	0	**NOTE:** Overrides Debug setting from `Common.Shared`. Additional parameters for debug/development use. This enables debugging modes for PM.
RmppDebug	0	**NOTE:** Overrides RmppDebug setting from `Common.Shared`. If 1, then log additional PM info with regards to RMPP (Reliable Message Passing Protocol).

Parameter	Default Value	Description
CS_LogMask	0x00000000	Alternative to use of `LogLevel`. For advanced users, these parameters can provide more precise control over per subsystem logging. For typical configurations, these should be omitted and the `LogLevel` parameter should be used instead. For each subsystem, there can be a `LogMask`. The mask selects severities of log messages to enable and is a sum of the following values: 0x1=fatal 0x2=actionable error 0x4=actionable warning 0x8=actionable notice 0x10=actionable info 0x20=error 0x40=warn 0x80=notice 0x100=progress 0x200=info 0x400=verbose 0x800=data 0x1000=debug1 0x2000=debug2 0x4000=debug3 0x8000=debug4 0x10000=func call 0x20000=func args 0x40000=func exit For embedded Fabric Manager, the corresponding Chassis Logging must also be enabled and SM configuration applies to all managers. For Host Fabric Manager, the Linux syslog service will need to have an appropriate level of logging enabled.
MAI_LogMask	0x000001ff	
CAL_LogMask	0x000001ff	
DVR_LogMask	0x000001ff	
IF3_LogMask	0x000001ff	
SM_LogMask	0x000001ff	
SA_LogMask	0x000001ff	
PM_LogMask	0x000001ff	
PA_LogMask	0x000001ff	
FE_LogMask	0x000001ff	
APP_LogMask	0x000001ff	

2.6. Subnet Manager Monitoring

This section provides information on the subnet manager monitoring and debug capabilities.

2.6.1. SM Logging and Debug

When nodes appear or disappear from the fabric, a message is logged. The SM Logging/Debug section defines how many messages are logged, and if additional SM sweep and SA query debug information is logged. The following is a sample of the SM Logging/Debug section of the Fabric Manager configuration file.

```
<!-- ***************** SM Logging/Debug ************************** -->
<!-- When nodes appear or disappear from the fabric, a message is logged -->
<!-- This can set a threshold on how many such messages to output per -->
<!-- sweep.  Once NodeAppearanceMsgThreshold messages are logged in a -->
<!-- given sweep, the remainder are output at a lower log level (INFO) -->
<!-- Hence avoiding excessive log messages when significant -->
<!-- fabric changes occur.  0 means no limit. -->
<NodeAppearanceMsgThreshold>100</NodeAppearanceMsgThreshold>
<SmPerfDebug>0</SmPerfDebug> <!-- log additional SM sweep info -->
<SaPerfDebug>0</SaPerfDebug> <!-- log additional SA query info -->
```

The parameters in the following table control SM Logging.

Table 11. SM Logging and Debug Parameters

Parameter	Default Value	Description
NodeAppearanceMsgThreshold	100	This can set a threshold on how many log messages to output per sweep. Once `NodeAppearanceMsgThreshold` messages are logged in a given sweep, the remainder are output at a lower log level (INFO) therefore avoiding excessive log messages when significant fabric changes occur. `0` means no limit.
SmPerfDebug	0	If `1`, then log additional SM sweep info.
SaPerfDebug	0	If `1`, then log additional SA query info. Log additional SA query info `SM_0_sa_debug_perf:dec` • `<Debug>0</Debug>` – #SM_0_debug:dec • `<RmppDebug>0</RmppDebug>` – #SM_0_sa_debug_rmpp:dec

2.6.2. SM Overrides of the Common.Shared Parameters

The `Common.Shared` parameters can be overridden in the SM using the parameters described in the following table.

Table 12. Additional SM Parameters

Parameter	Default Value	Description
Priority	0	0 to 15, higher wins.
ElevatedPriority		0 to 15, higher wins.

Parameter	Default Value	Description
LogLevel	2	**NOTE:** Overrides the `Common.Shared` LogLevel Settings. Sets log level option for SM: • 0 = disable the vast majority of logging output. • 1 = fatal, error, warn (syslog CRIT, ERR, WARN). • 2 = +notice, progress (syslog NOTICE, INFO). • 3 = +INFO (syslog DEBUG). • 4 = +VERBOSE and some packet data (syslog DEBUG). • 5 = +debug trace info (syslog DEBUG). This parameter is ignored for the Embedded Fabric Manager. Refer to the *CN5000 Commands Guide* for information on configuring chassis logging options.
LogFile		**NOTE:** Overrides `Common.Shared` settings. Sets log output location for SM. By default (or if this parameter is empty) log output is accomplished using syslog. However, if a LogFile is specified, logging will be done to the given file. LogMode further controls logging.
SyslogFacility	Local6	**NOTE:** Overrides `Common.Shared` settings. For the Host Fabric Manager, controls what syslog facility code is used for log messages. Allowed values are: auth, authpriv, cron, daemon, ftp, kern, local0-local7, lpr, mail, news, syslog, user, or uucp.
ConfigConsistencyCheckLevel	2	Controls the Configuration Consistency Check for SM. If specified for an individual instance of SM, overrides Shared settings. Checking can be completely disabled or can be set to take action by deactivating Standby SM if the configuration does not pass the consistency check criteria. • 0 = disable Configuration Consistency Checking • 1 = enable Configuration Consistency Checking without taking action (only log a message) • 2= enable Configuration Consistency Checking and take action (log message and move standby to inactive state)

2.6.3. LID

LID can be set for this SM as shown in the following table.

NOTE

If you change the default, and you are using multiple Fabric Managers in a fabric, it is recommended that you use the same SM LID setting on all Fabric Managers in the fabric.

Table 13. Additional SM LID Parameter

Parameter	Default Value	Description
LID	0	LID for this SM, 0=pick any available. 0 is recommended.

3. Maintenance

This chapter provides common procedures and information to maintain your CN5000 Omni-Path system.

3.1. Hardware Maintenance

3.1.1. Maintenance on an Active Cluster

This section describes the hardware maintenance that can be performed while the Omni-Path hardware is up and running.

IMPORTANT

Care must be taken as some of these activities may be disruptive while jobs are running.

3.1.1.1. Replacing CN5000 Switch Hot-Swappable Modules

3.1.1.1.1. Power Supply

To remove a CN5000 Switch power supply:

1. Press the thumb tab inward and grasp the handle to gently slide the module out of the slot.

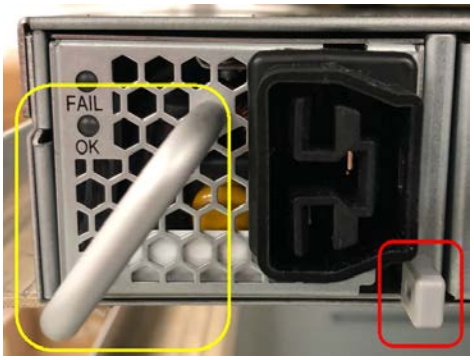

NOTE

Handle is highlighted yellow; thumb tab is highlighted red.

To insert a CN5000 Switch power supply:

1. Grasp the module by the handle and gently slide it into an open slot until the unit engages with the connector.

   When fully inserted, the module sits flush with the chassis and the thumb tab is locked.

3.1.1.1.2. Fans

To remove a CN5000 Switch fan:

1. Press the thumb tab inward and grasp the handle to gently slide the module out of the slot.

NOTE

Handle is highlighted yellow; thumb tab is highlighted red.

To insert a CN5000 Switch fan:

1. Grasp the module by the handle and gently slide it into an open slot until the module engages with the connector.

   When fully inserted, the module sits flush with the switch and the thumb tab is locked.

3.1.1.2. Replacing the DCS Modules

3.1.1.2.1. Leaf Modules

To remove a DCS Leaf module:

1. Grasp and pull the release arms (top and bottom) out and away from the module body.

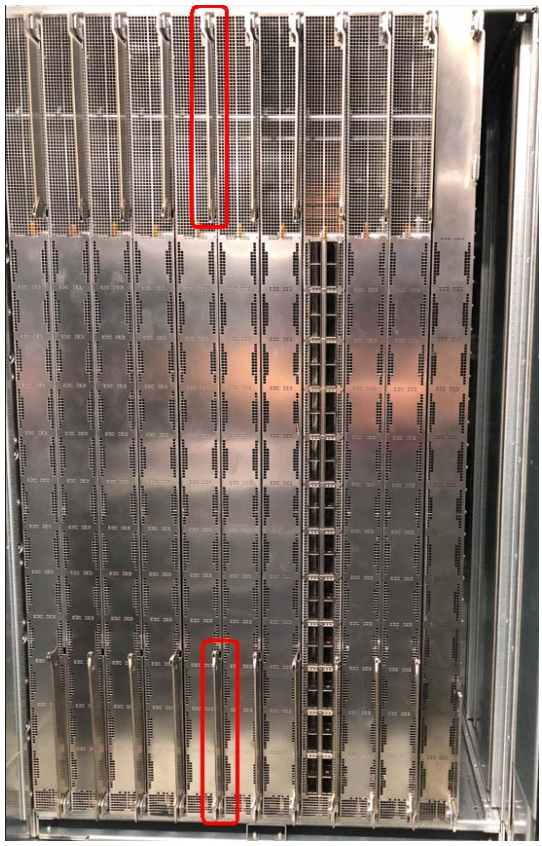

NOTE

Release arms are highlighted red.

2. Gently pull the module out of the chassis.

CAUTION

Never pull on the release arms when removing the modules from the slots.

To insert a DCS Leaf module into a chassis:

1. Grasp and pull the release arms (top and bottom) out and away from the module body.
2. Gently slide the module straight into an applicable open slot until it engages with the DCS chassis.

   CAUTION

   Never push on the release arms when sliding the modules into the applicable slots. Always apply pressure to the middle of the module.

3. Push the release arms in until they lock in place.

3.1.1.2.2. Fan Tray Modules

NOTE

Air baffles are located on the back side of the fan tray modules.

To remove a DCS Fan Tray module:

1. Grasp and squeeze the handles inward to release the module.

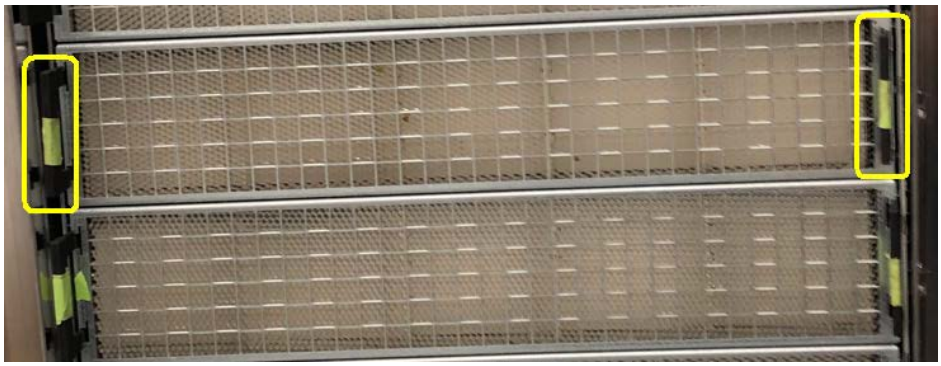

NOTE

Handles are highlighted yellow.

2. Gently slide the module out of the slot.

To insert a DCS Fan Tray module:

1. Gently slide the module into an open slot until the unit engages with the chassis.

   When fully inserted, the module locks into place and sits flush with the chassis.

3.1.1.2.3. Spine Modules

To remove the Spine module:

1. Remove the fan tray module as described in Section 3.1.1.2.2 “Fan Tray Modules”.

   Behind each fan tray are two Spine modules.
2. Press the thumb tabs inward to release the lock.

NOTE

Handles are highlighted yellow; thumb tabs are highlighted red.

3. Grasp the handles (if the top module) or the module itself (if without handles) to gently slide the module out of the slot.

To insert a Spine module:

1. Ensure the thumb tabs are pressed inward to retract the locking mechanism.
2. Gently slide the module straight into an applicable open slot until it engages with the DCS chassis.
3. Press the thumb tabs outward to engage the lock.
4. Reinsert the fan tray module as described in Section 3.1.1.2.2 “Fan Tray Modules”.

3.1.1.2.4. Management Modules

To remove a DCS Management module:

1. Grasp the bottom handle and lift upward to release the module

> **NOTE**
> Handle is highlighted yellow.

2. Gently pull the module out of the slot.

To insert a DCS Management module:

1. Gently slide the module straight into the open slot until it engages with the chassis. When fully inserted, the module is locked.

3.1.1.2.5. Power Supplies Modules

To remove a DCS Power Supply module:

1. Position the handle at the center of the fan, if previously moved off to the side.

2. Lift the release tab, then grasp the module by the handle and gently slide it out of the slot.

To install a DCS Power Supply module:

1. Grasp the module by the handle and gently slide it into an open slot until the unit engages with the connector.

   When fully inserted, the module sits flush with the chassis and the thumb tab is locked.

3.1.2. Offline Maintenance

This section describes the hardware maintenance that must be performed when the Omni-Path hardware is offline.

3.1.2.1. DCS Chassis Replacement

Prerequisites

- Ensure two or more people are available to lift the DCS when required.
- Ensure that all cables are labeled properly prior to removal.

Procedure

1. Power down the chassis using power button.

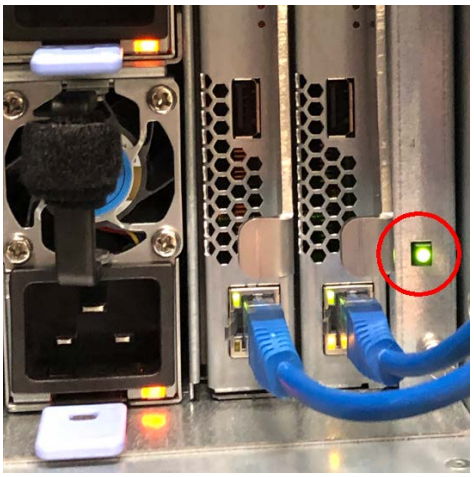

2. Disconnect all power and Ethernet wires.
3. Remove all ISL cables.
4. Remove all modules: leafs, spines, power supplies, and fans.

> **NOTE**
>
> Removing the modules reduces the weight that will make replacement easier.

5. Unscrew the chassis from the rails/racks.
6. Slide the chassis onto the lift platform.
7. Replace the chassis and insert the modules and power cables.
8. Check the firmware and update if needed.

9. Plug in all ISLs.

3.1.2.2. Switch Replacement

Prerequisites

- Ensure that all cables are labeled properly prior to removal.

Procedure

1. Disconnect the power and Ethernet cables.
2. Remove all cables.
3. Unscrew the CN5000 Switch from the rails/racks.
4. Slide the CN5000 Switch onto a lab cart.
5. Replace CN5000 Switch.
6. Insert the cables: Ethernet, Omni-Path, and power.
7. Check the firmware and update as needed.

3.1.2.3. Remove a Liquid-Cooled Switch

Prerequisites

Ensure the Switch has been shut down through the cluster management framework. Allow the Switch to reach a safe internal temperature (passive cool-down).

Procedure

1. Isolate and depressurize the cooling loop at the node.
   a. Close or isolate the branch of the cooling manifold feeding the node (if supported by rack/CDU design).
   b. Relieve local hydrostatic pressure at the quick-disconnect interface per service instructions.

   

   NOTE

   Ensure hoses are safely supported to avoid strain on the QDs.

2. Disconnect the quick-disconnect couplers by depressing the quick-disconnect release mechanism and separate the male/female halves. Switch-side fittings will automatically seal when disconnected.

3. To drain the Switch, remove both quick-disconnects from the switch. Alternatively, attach a hose to the quick-disconnects where the hose is open to the atmosphere.
4. Remove the hose(s) if necessary, and cap all liquid ports.
5. Remove the Switch (refer to Section 3.1.2.2 "Switch Replacement").
6. Store or prepare the Switch for transport per Cornelis packaging requirements.

NOTE

Package the switch to prevent physical damage. Insure for full replacement value.

3.1.2.4. SuperNIC Replacement

1. Power down the server.
2. Remove all cables.
3. Slide the server out.
4. Remove the cover from the server.
5. Remove the PCIe-Riser adapter from the PCIe slot, if applicable.
6. Remove the old SuperNIC from the slot.
7. Insert the new SuperNIC.
8. Reinsert the PCIe Riser adapter into PCIe slot.
9. Replace the cover on the server.
10. Push the server back into the rack.
11. Reinsert all cables.
12. Power on the server.

3.2. Software/Firmware Maintenance

This section describes the software and firmware maintenance including upgrades, updates, and installation of new modules.

3.2.1. Download the Firmware

Download the firmware or software using the following procedures.

1. Using a web browser, go to the Cornelis Customer Center.
2. Under Download Library, clear the navigation filters.
3. In the search box, enter your search string (for example, "firmware").

   The results are displayed.

4. Select one or more items and click **Download Selected**.
5. Review the Software License Agreement(s) and click **Accept** for each item.

   The firmware is saved to your computer.

3.2.2. Install the SuperNIC Firmware Update Tool

Prerequisites

- The OPX Software (containing updateAgent dependencies) has been installed on the target server.

Procedure

To install the SuperNIC Firmware Update Tool, perform the following steps.

1. Download and extract the Firmware Update TGZ package (contains `updateAgent`) from the Cornelis Customer Center.
2. Copy the `updateAgent` binary to the root home (/root) on the target server containing the SuperNIC you want to update.

3.2.3. Update the SuperNIC Firmware

Use the SuperNIC Firmware Update Tool to update your SuperNIC firmware.

Prerequisites

- The `updateAgent` must be copied to the server containing the SuperNIC you want to update.
- The hfi1 driver must be loaded before using the `updateAgent`.

  If not, you will see errors like "*Failed to open MAD port for HFI 0*" and "*Failed to build HFI list*".

Procedure

1. Obtain the SuperNIC firmware (`CN5000_SuperNICFirmware-<version>.pkg`) from the Cornelis Customer Center and copy the file to the server containing the SuperNIC you want to update.
2. Verify the hfi1 driver is loaded and working correctly.
   - `lsmod | grep hfi1` should return results
   - `opainfo` should have entries for all SuperNICs in the system

     If needed, load the hfi1 driver using `modprobe hfi1`, then recheck `opainfo`.

3. Check the current SuperNIC firmware version.

```
./updateAgent -V
HFI hfi1_0 activeComponentImageSetVersionString: <current version>
```

> **NOTE**
>
> You can use `./updateAgent -V -d all` to display all of the SuperNICs on a server.

4. If the Fabric Manager is running anywhere in the fabric, disable the active SuperNIC port.

```
opaportconfig disable -h1 -p2
```

Alternately, you can stop the Fabric Manager.

```
systemctl stop opafm
```

5. Update the SuperNIC.

```
./updateAgent /path/to/firmware.pkg
```

> **NOTE**
>
> - To update all SuperNICs on a server, you can use:
>
>   ```
>   ./updateAgent -d all /path/to/firmware.pkg
>   ```
>
> - To update all SuperNICs in a fabric, you can use tools such a pdsh command as shown in the following example:
>
>   ```
>   pdsh -w <hostfile> updateAgent -d all /path/to/firmware.pkg
>   ```

6. Check the current SuperNIC firmware version and verify that the new version status is `pendingComponentImageSetVersionString`.

```
./updateAgent -V
HFI hfi1_0 activeComponentImageSetVersionString: <old version>
HFI hfi1_0 pendingComponentImageSetVersionString: <new version>
```

7. Power cycle the server.
8. Check the current SuperNIC firmware version again and verify that the status is `activeComponentImageSetVersionString`.

```
./updateAgent -V
HFI hfi1_0 activeComponentImageSetVersionString: <new version>
```

If the firmware is still pending, power cycle the server using BMC.

9. If you stopped the Fabric Manager in step 4, restart it.

```
systemctl start opafm
```

3.2.4. Update the Switch Firmware

If you are updating both BMC and ASIC firmware, you must update the BMC firmware first.

NOTE

- In the following instructions for the Pull Method, `user@hostname` implies DNS is configured on the switch. If using static IP addresses, replace this text with the IP address of the switch containing the .pkg file.
- When using `firmware update` at the switch CLI to transfer ("pull") firmware from a remote server, you'll be prompted for the remote server password.
- When using SCP on a remote server to transfer ("push") firmware to the switch, you'll be prompted for the switch password.
- It may take up to 20 minutes (for CN5000 Switch) or 40 minutes (for DCS) for a firmware update to complete.

Perform the following steps to update your switch firmware.

NOTE

If you are updating multiple switches, repeat these steps for each switch.

1. Download and extract the target Switch firmware package files (BMC Firmware and/or Switch Firmware) from the Cornelis Customer Center onto a server on the same Ethernet network as the switch to be updated.

Updating the BMC Firmware

NOTE

During the initial installation of this new version (after the forced reboot) or when inserting new boards, the BMCs may require multiple updates and reboots (up to 2) to synchronize the firmware across the entire DCS. After this initial synchronization, future updates should only require a single final reboot to apply. As long as the ASIC is off, this process should be automatic.

2. Run the `firmware update` command to begin the update process.
   - **Pull Method:** If not already logged in, log into the switch using the admin account. Specify the `user@hostname:/path/to/file.pkg` path. Enter the password of the host when prompted.

```
admin@CNEdge -> firmware update user@hostname:/path/to/CN5000_BMCFirmware-
<version>.pkg
root@hostname's password:
Copying firmware image to staging area...
Firmware update started. Wait (up to 20 minutes), check status with "firmware
update -s", and initiate a reboot when ready
```

   - **Push Method:** Specify the `admin@switchName:/tmp/images` destination path. Enter the switch password if/when prompted.

```
[user@servername ~]# scp -O root/fw/CN5000_BMCFirmware-<version>.pkg
admin@switchname:/tmp/images
admin@switchname's password:
Copying firmware image to staging area...
Firmware update started. Wait (up to 20 minutes), check status with "firmware
update -s", and initiate a ‘reboot force’ when ready
```

Updating the Switch Firmware (ASIC)

3. Run the `firmware update` command to begin the update process.
   - **Pull Method:** If not already logged in, log into the switch using the admin account. Specify the `user@hostname:/path/to/file.pkg` path. Enter the password of the host when prompted.

```
admin@CNEdge -> firmware update user@hostname:/path/to/CN5000_SwitchFirmware-
<version>.pkg
root@hostname's password:
Copying firmware image to staging area...
Firmware update started. Wait (up to 20 minutes), check status with "firmware
update -s", and initiate a 'reboot -f' when ready
```

   - **Push Method:** Specify the `admin@switchName:/tmp/images` destination path. Enter the switch password if/when prompted.

```
[user@servername ~]# scp -O root/fw/CN5000_SwitchFirmware-<version>.pkg
admin@switchname:/tmp/images
admin@switchname's password:
Copying firmware image to staging area...
Firmware update started. Wait (up to 20 minutes), check status with "firmware
update -s", and initiate a ‘reboot force’ when ready
```

Completing the Update

4. Check the status of the update using `firmware update -s`.

```
admin@CNEdge -> firmware update -s
BMC:
    Image 1: Booted and Active
    Image 2: Currently updating
ASIC A:
    Image 1: Booted and Active
    Image 2: Currently updating
```

   changes to

```
admin@CNEdge -> firmware update -s
BMC:
    Image 1: Booted and Active
    Image 2: Staged for update
ASIC A:
    Image 1: Booted and Active
    Image 2: Staged for update
```

   When the firmware shows `Staged for update`, the switch is ready for reboot.

5. Reboot the switch.

```
admin@CNEdge -> reboot force
Rebooting in 1 second(s)Lost Communication with server
Connection to <hostname> closed by remote host.
Connection to <hostname> closed.
```

> **NOTE**
>
> If you try to reboot BEFORE the firmware is in `Staged for update`, you will receive the following error:
>
> ```
> Error during firmware update, rebooting now could be dangerous. Are you sure you wish to continue?
> ```
>
> Type **no** and wait for the status to change.

After updating the switch firmware, check the versions.

```
admin@CNEdge -> firmware version
Firmware Versions:

Switch BMC Chip version: <current version>
ASIC Chip A version: <current version>
```

3.2.5. Update the OPX Software

This section provides information and procedures to update to the OPX Software.

Prerequisites

Prior to updating the OPX Software, ensure the following items have been completed:

- Review the Release Notes for a list of compatible software.
- Uninstall all versions of third-party IB stacks.
- Back up the following configuration files, if applicable, in case the upgrade fails:
  - `/etc/opa-fm/opafm.xml`
  - `/etc/opa/*`
  - `/etc/sysconfig/opa/*`
  - `/var/usr/lib/opa/analysis/baseline/*`

NOTE

Refer to the OS documentation for a list of any other OS-specific files that should be included in any backups.

Update Sequence

Optimally, to update the OPX Software, you perform the following sequence:

1. Stop all standby Fabric Managers.
2. Stop the primary Fabric Manager.
3. Upgrade the nodes.
4. Reboot the nodes (if autostart was not enabled).
5. Update the `opafm.xml` file, as needed.
6. Start the Fabric Manager.

3.2.5.1. Stopping the Fabric Manager

To stop the Fabric Manager, perform the following steps:

1. Log into the Fabric Manager system as `root` or as a user with `root` privileges.
2. Stop the Fabric Manager:

```
systemctl stop opafm
```

3.2.5.2. Upgrading the OPX Software

To upgrade the CN5000 OPX Software, you first upgrade the management node then the remaining servers.

Assumption

- You have downloaded the required software RPMs for upgrade.
- You have stopped all standby FMs, followed by the primary FM.
- You are logged in to the target Management Node.

Procedure

The following instructions provide the steps for upgrading the host software on the first management node and then the remaining nodes.

NOTE

If you plan to use a local repository and the TGZ package from the Cornelis Customer Center, download and extract the package to the management node.

Download and install the OPX Software package:

1. Create a repository folder or confirm that the following already exists.
   - For RHEL:

     ```
     /etc/yum.repos.d/
     ```

   - For SLES:

     ```
     /etc/zypp/repos.d/
     ```

   - For Ubuntu:

     ```
     /etc/apt/sources.list.d/
     ```

2. Create a repository file.
   - For RHEL:

```
sudo vi /etc/yum.repos.d/cornelis-repository.repo
[Cornelis-Package]
name=Cornelis Repository
baseurl=file:///<file path to local downloaded updated tar file>
enabled=1
gpgcheck=1
```

   - For SLES:

```
sudo vi /etc/zypp/repos.d/cornelis-repository.repo
[Cornelis-Package]
name=Cornelis Repository
baseurl=file:///<file path to local downloaded updated tar file>
autorefresh=1
enabled=1
gpgcheck=1
```

   - For Ubuntu:

```
sudo vi /etc/apt/sources.list.d/cornelis-repository.list
deb [trusted=yes] file:<file path to local downloaded updated tar file>/ ./
```

3. Import the GPG Key.
   - For RHEL and SLES:

```
rpm --import <file path to local downloaded updated tar file>/CN5000-Packages-
Public-GPG-Key.asc
```

   - For Ubuntu:

```
sudo gpg --dearmor -o /etc/apt/trusted.gpg.d/cornelis.gpg <file path to local
downloaded updated tar file>/CN5000-Packages-Public-GPG-Key.asc
```

4. Clean and update the package Index.
   - For RHEL:

```
sudo dnf clean all
sudo dnf makecache
```

   - For SLES:

```
sudo zypper clean
sudo zypper refresh
```

- For Ubuntu:

```
sudo apt clean
sudo apt update
```

5. Determine the meta package you need to use for installation.

NOTE

Ubuntu uses dashes in place of underscores in the following list.

- cn5000_pkgs_non_gpu_meta - CPU-based installation
- cn5000_pkgs_cuda_meta - NVIDIA GPU installation
- cn5000_pkgs_rocm_meta - AMD GPU installation

CAUTION

Installing meta packages will override existing packages. If multiple versions are desired (for example, GPU/non-GPU), additional components can be installed directly using the package manager, but not using the meta packages.

6. Upgrade the package using one of the following commands.
   - For RHEL:

```
sudo dnf install <META_PACKAGE_NAME>
```

   - For SLES:

```
sudo zypper install <META_PACKAGE_NAME>
```

NOTE

If zypper prompts for confirmation, you can select the default to continue.

- For Ubuntu:

```
sudo apt install <META_PACKAGE_NAME>
```

NOTE

If you receive an error related to required dependencies that are not the latest versions available on Ubuntu, Cornelis recommends that you install and use `aptitude` in place of `apt` in the command.

Alternately, you can use `apt` or `apt-get` with `--mark-auto` to manually install the dependencies.

7. Perform the following:
   - For RHEL and SLES, continue to the next step to download, build, and install the kmod-opxs-kernel-updates packages.
   - For Ubuntu, reboot the server to finalize the installation, then go to Step 13 to install the software to each remaining server.

NOTE

The opxs-modules-dkms packages are already installed with the Ubuntu meta package.

Download, build, and install the kmod-opxs-kernel-updates packages (RHEL and SLES only):

8. Copy the relevant RPM source file (src.rpm) to a local directory.

IMPORTANT

- Ensure you are using the correct src.rpm file to build and install the version of the opx-kernel-updates relevant to your hardware configuration (AMD GPU, NVIDIA GPU, or CPU-only). Three opxs-kernel-updates RPM source files are provided, one each for AMD GPU, NVIDIA GPU, and CPU-only.
- RHEL Only: To rebuild the source RPM, you must install `kernel-rpm-macros` (for AMD GPUs) and `kernel-abi-stablelists`.
- SLES Only: You need to build `unifdef` before running rpmbuild:

```
cd /lib/modules/$(uname -r)/source/scripts
gcc -o unifdef unifdef.c
```

9. Build the source RPM.

```
rpmbuild --rebuild --define "_topdir <local_directory>" --define 'dist %{nil}' --
target x86_64 --define 'kver $(uname -r)' <src.rpm_file>
```

10. Verify the RPMs are built by changing directories to the `<local_directory>/RPMS/x86_64` directory and running `ls -1`.

    Output should be:

```
kmod-opxs-kernel-updates-<build_version>.x86_64.rpm
opxs-kernel-updates-devel-<build_version>.x86_64.rpm
```

11. Upgrade the new RPMs.
    - For RHEL:

```
sudo dnf install kmod-opxs-kernel-updates-<build_version>.x86_64.rpm
opxs-kernel-updates-devel-<build_version>.x86_64.rpm
```

    - For SLES:

```
zypper --no-gpg-checks install opxs-kernel-updates-devel-
<build_version>.x86_64.rpm
opxs-kernel-updates-kmp-default-<build_version>.x86_64.rpm
```

12. Reboot the primary node to apply changes.

Install the OPX Software to each remaining server:

13. Copy the Repository folder, the tar file (if using a local repository), and the `<local_directory>` containing the newly built RPMs, to each remaining server.

> **IMPORTANT**
>
> The folders and files must be in the same directories as was in the management node.

14. Extract the TGZ package.
15. Upgrade the software on the remaining nodes using the commands in steps 6 and 11.
16. Reboot the rest of the nodes to apply changes.

17. After update, disable the repository on every node using one of the following methods:
    - For RHEL, run the following commands:

```
sudo dnf config-manager --set-disabled <repo>
```

    - For RHEL and SLES, in the repository file, change `enable` to `0`.
    - For Ubuntu, in the repository (`.list`) file, comment out the following line:

```
#deb [trusted=yes] file:/<file path to local downloaded updated tar file>/ ./
```

3.2.5.3. Update the Fabric Manager opafm Configuration File

The `opafm` configuration file is designed so that any new parameters that might not appear in your old `opafm.xml` file will automatically be set to their defaults in the software. If you have customized your pre-upgrade `opafm.xml` file, it is recommended that you manually update your new `opafm.xml` file with the latest comments and defaults. This will aid you with future customizations and make comparisons with the `opafm.xml` in future upgrades easier.

IMPORTANT

The `opafm` configuration file **must** be the same on all management nodes.

Assumption

- You have upgraded the Fabric Manager, accepting all the defaults.

Procedures

Perform the following steps to transfer customizations into a new `opafm.xml` file:

1. Create a new `opafm.xml` using the `/usr/share/opa-fm/opafm.xml` file as a base for the new file by copying this read-only file to a temporary location to make the subsequent edits.
2. Update the new `opafm.xml` with the required customizations.
3. Use the new file from the previous step to replace the `/etc/opa-fm/opafm.xml` file.
4. Restart the Fabric Manager.
    - If you previously upgraded the management node, type `reboot` and press **Enter**.

NOTE

If the Omni-Path Autostart was disabled, restart the primary Fabric Manager, followed by the standby Fabric Manager. Refer to Section 3.2.5.4 “Starting the Fabric Manager”.

- If you only updated the `opafm.xml` file, type `systemctl restart opafm` and press **Enter**.

5. Copy the `opafm.xml` file to the standby Fabric Managers, replacing the `/etc/opa-fm/opafm.xml` file.
6. Restart the standby Fabric Managers.

3.2.5.4. Starting the Fabric Manager

This section provides instructions for starting the Fabric Manager.

IMPORTANT

The Fabric Manager is always run from a node connected to the latest generation switch (for example, a CN5000 switch).

Prerequisites

Prior to starting the Fabric Manager, you must edit the `opafm.xml` file (located under `/etc/opa-fm/`) and ensure the following:

- HFI 1, port 2 is enabled for fm0:

```
<Start>1</Start>  <-- 1 signifies enabled
<Name>fm0</Name>
<Hfi>1</Hfi>
<Port>2</Port>
```

NOTE

Also note that a single port SuperNIC appears as dual port in `opainfo` output.

- CableInfoPolicy is set to None:

```
<CableInfoPolicy>None</CableInfoPolicy>
```

Procedure

To start the Fabric Manager, perform the following steps:

1. Log into the Fabric Manager system as `root` or as a user with `root` privileges.
2. Start the Fabric Manager:

```
systemctl start opafm
```

3. Verify that all the tasks are up and running:

```
systemctl status opafm
```

NOTE

The default configuration runs the Fabric Manager on port 2 of the first SuperNIC in the system.

3.2.6. Remove OPX Software from a Host

If you need to remove the OPX Software from a host due to an issue with an OS upgrade, perform the following steps.

1. Uninstall the software using the command.
   - Using RHEL:

     ```
     sudo dnf remove <META_PACKAGE_NAME>
     ```

   - Using SLES:

     ```
     sudo zypper remove <META_PACKAGE_NAME>
     ```

   - Using Ubuntu:

     ```
     sudo apt remove <META_PACKAGE_NAME>
     ```

2. Remove kernel packages.
   - Using RHEL:

     ```
     sudo dnf remove kmod-opxs-kernel-updates opxs-kernel-updates-devel
     ```

   - Using SLES:

     ```
     sudo zypper remove kmod-opxs-kernel-updates opxs-kernel-updates-devel
     ```

   - Using Ubuntu:

     ```
     sudo apt remove opxs-modules-dkms*
     ```

3. Remove residual packages.
   - Using RHEL:

     ```
     sudo dnf remove opa-*
     sudo dnf remove libfabric-*
     sudo dnf remove libpsm2-*
     ```

   - Using SLES:

     ```
     sudo zypper remove opa-*
     sudo zypper remove libfabric-*
     sudo zypper remove libpsm2-*
     ```

   - Using Ubuntu:

```
sudo apt remove opa-*
sudo apt remove libfabric-*
```

4. Reboot.

4. Troubleshooting

For fabric monitoring tools and commands used to assist with troubleshooting, refer to Section 2.3 "Common Fabric Monitoring Tools".

4.1. Hardware Troubleshooting

4.1.1. Identifying Issues from LEDs

LEDs offer quick visual diagnostics, using color and flashing patterns to indicate power, network, drive activity, and hardware errors, enabling rapid troubleshooting and reducing downtime.

Refer to the *CN5000 Product Family Description Guide*, "Fabric Hardware Components" for the location and meaning of each product/component LED.

4.1.2. Identifying Suspect Cables

Often the first indication of possible signal integrity and link issues associated with suspect cables appear in the Fabric Manager log messages. You can observe fabric issues through several methods including log messages in the syslog and monitoring tools.

4.1.2.1. Common PM Log Messages

This section provides some of common PM log messages, what they mean and what they could indicate.

4.1.2.1.1. General Integrity Threshold

As the PM sweeps more frequently, the PM is often the first to log messages indicating signal integrity issues.

Format:

```
fm0_sm[<PID>]: WARN [PmEngine]: PM: Integrity of <VALUE> Exceeded Threshold of 100.
<NODE> Guid <GUID> LID <LID> Port <#>  Neighbor: <NEIGHBOR_NODE> Guid <NEIGHBOR_GUID> LID
<NEIGHBOR_LID> Port <#>
fm0_sm[<PID>]: WARN [PmEngine]: PM: <COUNTER_DATA>
```

Example:

```
fm0_sm[12345]: WARN [PmEngine]: PM: Integrity of 600 Exceeded Threshold of 100. compute001
hfi1_0 Guid 0x0011750101010101 LID 0x3 Port 1  Neighbor: edge003 Guid 0x00117501cafebeef
LID 0x2 Port 2
fm0_sm[12345]: WARN [PmEngine]: PM: LQI=1
```

NOTE

Only the first Pm.ThresholdsExceededMsgLimit.integrity messages (default of 10) will print each sweep as WARN. The remaining messages will print as INFO (which, by default, will not print on most rsyslog configs).

4.1.2.1.2. Port Bounce Integrity Threshold

In the event that a port is not accessible by the PM, often the neighbor may report an integrity message similar to the General Integrity Threshold Message that the port/link is in the DOWN state (`LQI=0`).

When rebooting a node or during maintenance, these messages can be quite common. The PM often queries ports before the SM finalizes its topology. If the messages do not repeat after the reboot or maintenance, they can be ignored.

4.1.2.1.3. Query Failures

If the PM is unable to query the counters on a port, it reports an error. This error can be normal if the server or switch is rebooting at the time.

Occasionally a failure occurs when a port fails to respond because of an issue (such as signal integrity) along the request or response route.

Format:

```
fm0_sm[<PID>]: WARN [PmEngine]: PM: PmPrintFailPort: Unable to Get(PortStatus) <NODE> Guid
<GUID> LID <LID> Port <#>
```

Example:

```
fm0_sm[12345]: WARN [PmAsyncRcv]: PM: PmPrintFailPort: Unable to Get(PortStatus)
compute004 hfi1_0 Guid 0x0011750101010102 LID 0x1f Port 1
```

NOTE

Only the first ten (Pm.SweepErrorsLogThreshold) Sweep Errors messages will print each sweep as WARN. The remaining messages will print as INFO.

PM Sweep Failures occur any time the PM fails to query the counters on a port during a sweep. At the end of any PM sweep that encountered errors, the PM provides a summary of the number of nodes and ports that it was unable to access as shown in the example below.

Format:

```
fm0_sm[<PID>]: WARN [PmEngine]: PM: PmSweepAllPortCounters: Unable to get <#> Ports on <#>
Nodes
```

Example:

```
fm0_sm[12345]: WARN [PmEngine]: PM: PmSweepAllPortCounters: Unable to get 277 Ports on 60
Nodes
```

4.1.2.2. Common SM Log Messages

This section provides some of the common SM log messages, what they mean, and what they could indicate.

4.1.2.2.1. Node Appearance and Disappearance Messages

An Appearance Message indicates that a node has entered the fabric topology during the previous sweep. Conversely, a Disappearance Message indicates that a node has left the fabric topology during the previous sweep.

When a node leaves the fabric, the disappearance message is often accompanied by additional sweep failure messages.

The disappearance of a node in one sweep followed by an appearance of the same node in the next sweep usually indicates *intermittent issues* where the node fails to respond in time to the SM and is marked as no longer in the fabric topology until it can successfully respond.

Disappearance Format:

```
fm0_sm[<PID>]: <FM_NODE>; MSG:NOTICE|SM:<FM_NODE>:port <#>
|COND:#4 Disappearance from fabric
|NODE:<NODE>:port <#>:<GUID>
|LINKEDTO:<NEIGHBOR_NODE>:port <#>:<NEIGHBOR_GUID>
|DETAIL:Node type: <NODETYPE>
```

Disappearance Example:

```
fm0_sm[12345]: fm001; MSG:NOTICE|SM:fm001:port 1
|COND:#4 Disappearance from fabric
|NODE:compute002 hfi1:port 1:0x00117501deadbeef
|LINKEDTO:edge002:port 7:0x00117501deadcafe
|DETAIL:Node type: hfi
```

Appearance Format:

```
fm0_sm[<PID>]: <FM_NODE>; MSG:NOTICE|SM:<FM_NODE>:port <#>
|COND:#3 Appearance in fabric
|NODE:<NODE>:port <#>:<GUID>
|LINKEDTO:<NEIGHBOR_NODE>:port <#>:<NEIGHBOR_GUID>
|DETAIL:Node type: <NODETYPE>
```

Appearance Example:

```
fm0_sm[12345]: fm001; MSG:NOTICE|SM:fm001:port 1
|COND:#3 Appearance in fabric
|NODE:compute002 hfi1:port 1:0x00117501deadbeef
```

```
|LINKEDTO:edge002:port 7:0x00117501deadcafe
|DETAIL:Node type: hfi
```

4.1.2.2.2. Discovery Failures

Discovery failure messages are often seen during the beginning of the SM sweep where the FM runs through a quick discovery process to build a topology of the fabric for later steps in the SM sweep.

The message shown in the example below is usually accompanied by additional error messages. The node indicated in the failure is often the upstream port, which is the port the FM can see on the link (usually the switch port on a SuperNIC-to-switch link).

Format:

```
fm0_sm[<PID>]: WARN [topology]: SM: topology_discovery: unable to setup port[<#>] of node
<NAME>, nodeGuid <GUID>, ignoring port!
```

Example:

```
fm0_sm[12345]: WARN [topology]: SM: topology_discovery: unable to setup port[15] of node
edge002, nodeGuid 0x00117501deadcafe, ignoring port!
```

4.1.2.2.3. Setup Node Failures

A Setup Node Failure is the most common failure to accompany the Discovery Failure message. The failure occurs when the initial packet is sent across the wire and the SM did not receive a response (packet) and timed out. A status code of *7* is the common indicator for a timeout.

Format:

```
fm0_sm[<PID>]: WARN [topology]: SM: sm_setup_node: Get NodeInfo failed for nodeGuid <GUID>
port <#>, via node <NEIGHBOR_NODE> nodeGuid <NEIGHBOR_GUID> port <#>; status=7
```

Example:

```
fm0_sm[12345]: WARN [topology]: SM: sm_setup_node: Get NodeInfo failed for nodeGuid
0x00117501deadbeef port 1, via node edge002 nodeGuid 0x00117501deadcafe port 7; status=7
```

4.1.2.2.4. Programming Failures

Programming Failures (such as SCVL_t/nt) can occur intermittently when a node fails to respond during one of the programming phases of the sweep.

In the example shown below, a failure occurred when attempting to program a port. The response could not be completed successfully and the port was *marked down* indicating that the port was not part of the fabric topology.

Format:

```
fm0_sm[<PID>]: WARN [topology]: SM: sm_initialize_Switch_SCVLMaps: Failed to set SCVL_t
Map for node <NODE> nodeGuid <GUID> output port <#>
```

Example:

```
fm0_sm[12345]: WARN [topology]: SM: sm_initialize_Switch_SCVLMaps: Failed to set SCVL_t
Map for node edge003 nodeGuid 0x00117501cafebeef output port 13
```

4.1.3. Diagnosing Bad Cables

Once a suspect cable is identified, you can determine if it is *bad* and needs to be replaced by using several automated and manual processes.

4.1.3.1. Intermittent Link Quality Issues and Port Bounces

The Link Quality Indicator (LQI) provides a simple way to identify links with poor signal integrity. Links with poor LQI will be visible in PM logs, `opatop`, `opainfo`, `opalinkanalysis`, `opareport`, and other tools. Be aware that low LQI can occur when links are intentionally going down, such as during a device reboot or cable maintenance (replacement, reseating, and more). For more information on LQI, refer to Section 4.2.5.2.1 "Link Quality Indicator (LQI)".

Intermittent Link Quality issues often lead to random port bounces when the link exceeds a Bit Error Threshold and attempts to retrain the link. This may indicate that a port or cable is in the process of failing. You can perform the Section 4.1.3.4 "Cable Swap Test" to determine quickly if the issue is the port or the cable.

Additionally, sometimes if the failure occurs coinciding with any physical maintenance in the area (or at the port), you can reseat the cable, card, or blade to fix the issue.

4.1.3.2. High Error Counts

Once a link is identified as a possible *bad* link, you can observe the performance counters while the ports are in various states of load to determine the issue. For instance, a port with a significantly high rate of errors may point to a link that needs to be replaced. For a description of various port counters, refer to Section 4.2.5 "Port Counters".

4.1.3.3. Slow Links (LinkWidthDownGrade)

Slow Links are links that are not operating at the greatest supported bandwidth. For Omni-Path 400G, normal link operation would be 4x lane width with a lane speed of 100G.

Slow links usually indicate that a LinkWidthDowngrade event has occurred. One or more of the lanes experienced enough link issues that they could not continue to run. In order to avoid bouncing the whole link, the affected lanes were dropped.

4.1.3.4. Cable Swap Test

The Cable Swap Test is a simple way to determine if a link problem occurs with the cable.

Change out the suspect cable with a known, working cable and observe the results:

- If the problem goes away, then the old cable was bad.
- If the problem does not go away, then the problem resides elsewhere, such as in the port hardware on either end of the link.

4.1.3.5. Port Issues

When a link problem occurs with the ports, you need to determine which side is causing the issue. Switching one port to a new node can help to identify between two ports.

- If the issue is on the switch side, you can move to an unused port or swap out the switch (or Director Class leaf or spine). Reseating the leaf or spine may also help.
- If the issue is on the SuperNIC port, you can reseat the SuperNIC card to fix the issue. However, if the issue persists, then the SuperNIC card may need to be investigated. It is also possible that the issue may be in the server itself. If so, you can move the SuperNIC card to another slot or server to see if the issue is resolved.
- If the issue did not move to the new ports or stay with the old ports, then it is possible the issue was fixed by reseating the cable(s). Several repetitions might be needed to verify.

4.1.4. Mitigating Link Issues

For certain Fabric Manager configurations, you can allow for either more reliability or better performance during production runs.

The following sections describe the cost and benefits of changing certain Fabric Manager configuration options.

4.1.4.1. LinkPolicy

Several Link Policy options are available for both SuperNIC and ISL links. For each link type, the following options are included:

- MaxDroppedLanes sets the number of lanes on a link that the port can downgrade before bouncing itself and bringing up the link again. This value alters the LinkWidthDowngradeEnabled field on the port's portinfo.

  Changing the value of MaxDroppedLanes to 0 will prevent a link from ever operating in a downgraded mode. Instead, it will program an option on the port that will trigger the port to bounce itself instead of downgrading.

- WidthPolicy prevents the link activation of ports when the minimum configured link width is not met. Do not change this value without significant knowledge of what it will do.

IMPORTANT

When forcing a link policy change using the switch CLI, you will need to bounce the link from the switch side.

4.1.4.2. Timeouts

Several timeout mechanisms are available in the SM; some are replicated in the PM. This section describes the configuration options for *per packet timeouts* and *cumulative sweep timeouts*, including what effects will occur when you change them.

4.1.4.2.1. Per Packet Timeouts

Three main options are used for dealing with Per Packet Timeouts in both the SM and the PM:

- MaxAttempts
- RespTimeout
- MinRespTimeout

The following two modes use the options above to manage timeouts:

- The Increasing Timeouts mode is the default mode. When a packet is sent, it will wait up to the non-zero MinRespTimeout value (35 ms). If it times out, the packet will continue to retry at increasing multiples of MinRespTimeout until the timeout exceeds MaxAttempts multiplied by RespTimeout. For example, the first timeout is 35 ms, the second will be 70 ms, and so on.
- The Exact Timeout mode requires MinRespTimeout to be set to zero. When this is set, each packet will wait for RespTimeout and retry MaxAttempts.

4.1.4.2.2. Cumulative Sweep Timeouts

Cumulative Sweep Timeouts using the CumulativeTimeoutLimit option are available only in the SM. It is used to prevent too many nodes from timing out in one sweep, potentially extending a sweep significantly. The CumulativeTimeoutLimit value is in seconds and, once exceeded, will start to skip nodes instead of retrying. This value will allow for longer timeouts without also increasing the sweep time when multiple ports are experiencing errors.

4.1.4.3. Adaptive Routing LostRoutesOnly

When Adaptive Routing is enabled, rerouting for a port going down is always enabled. LostRouteOnly disables the congestion-based rerouting. When a port goes down, to prevent packets from being dropped, the switch will automatically reroute traffic when it finds a different route. For information about how to set up adaptive routing, refer to *CN5000 Topologies and Routing Guide*, "Adaptive Routing".

When LostRouteOnly is disabled, the switch also reroutes traffic when a specific congestion threshold is reached.

Though setting LostRouteOnly may not ease congestion, it will allow failures to be more optimally handled.

4.1.4.4. PortErrorAction

PortErrorAction is an SM configuration option that is programmed on the ports, directing them to automatically bounce if certain errors occur. These errors are listed in the configuration file.

The benefit of this option is that some issues can only be fixed by bouncing the port; if not bounced, the port may hang which could block traffic until fixed. Conversely, bouncing the port will interrupt traffic; and, if the issue was not causing the port to stall, bouncing may have been unnecessary.

4.1.4.5. Ports Speed

Switch ports may appear offline or down when an incorrect port speed is set on the link. Ensure that the attached cable supports the speed by using the Switch CLI command `hardware cable`.

Example output:

```
root@CNEdge:~# hardware cable
Port    Cable       Vendor          Cable  Part            Rv Serial          Spd sup  Temp
Name    Type        Name            Length Number             Number          25G 100G (C)
------- ----------- --------------- ------ --------------- -- --------------- --- ---- -----
1A      QSFP ACC    Amphenol          5.0m NJAAL6-CN05     A  APF2442N052U17   Y   Y     0.0
3A      QSFP ACC    Amphenol          3.0m NJAALR-CN03     A  APF2445N034R1L   Y   Y     0.0
5A      QSFP ACC    Amphenol          3.0m NJYKLR-CN03     A  APF2446N0374E7   Y   Y     0.0
6A      QSFP ACC    Amphenol          3.0m NJYKLR-CN03     A  APF2446N0374E7   Y   Y     0.0
7A      QSFP AOC    Hisense           7.0m DQF8503-4C07    02 S5189H700TR      Y   N    38.0
```

```
8A      QSFP AOC    Hisense            7.0m DQF8503-4C07     02 S5189H7005V       Y   N    35.0
9A      QSFP AOC    Hisense            7.0m DQF8503-4C07     02 S5189H7006H       Y   N    35.0
11A     QSFP ACC    Amphenol           4.0m NJAAL1-CN04      A  APF2520N04W98A    Y   Y     0.0
12A     QSFP ACC    Amphenol           5.0m NJAAL1-CN05      A  APF2509N054G63    Y   Y     0.0
13A     QSFP ACC    Amphenol           5.0m NJARL1-CN05      A  APF2509N054G9R    Y   Y     0.0
14A     QSFP DAC    Amphenol           2.0m NJAAK3-CN02      A  APF2509N024G4B    Y   Y
15A     QSFP ACC    Amphenol           5.0m NJAAL6-CN05      A  APF2442N052U1J    Y   Y     0.0
16A     QSFP AOC    FINISAR CORP.     10.0m FCBR4X0QE1C10COR A0 YV936KD           N   Y    48.5
17A     QSFP ACC    Amphenol           4.0m NJAAL6-0004      A  APF23430047K2Y    Y   Y     0.0
18A     QSFP AOC    FINISAR CORP.     10.0m FCBR4X0QE1C10    A0 YV8GRZ0           N   Y    48.4
19A     QSFP AOC    Hisense            7.0m DQF8503-4C07     02 S5189H7003K       Y   N    38.0
21A     QSFP AOC    Hisense            7.0m DQF8503-4C07     02 S5189H7002M       Y   N    37.0
22A     QSFP ACC    Amphenol           5.0m NJAAL6-CN05      A  APF2442N052U16    Y   Y     0.0
24A     QSFP ACC    Amphenol           5.0m NJAAL1-CN05      A  APF2509N054G66    Y   Y     0.0
25A     QSFP DAC    Amphenol           2.0m NJAAK3-CN02      A  APF2509N024G4C    Y   Y
26A     QSFP AOC    Hisense            7.0m DQF8503-4C07     02 S5189H70081       Y   N    39.0
27A     QSFP AOC    Hisense            7.0m DQF8503-4C07     02 S5189H700AJ       Y   N    39.0
28A     QSFP AOC    Hisense            7.0m DQF8503-4C07     02 S5189H700P0       Y   N    35.0
29A     QSFP AOC    Hisense            7.0m DQF8503-4C07     02 S5189H7004Q       Y   N    39.0
30A     QSFP AOC    Hisense            7.0m DQF8503-4C07     02 S5189H7002C       Y   N    39.0
31A     QSFP AOC    Hisense            7.0m DQF8503-4C07     02 S5189H7005X       Y   N    30.0
32A     QSFP AOC    Hisense           15.0m DQF8503-4C15     02 S5195AF000M       Y   N    35.0
33A     QSFP DAC    FCI Electronics    1.0m 10131941-2010LF  2  CN1515FA102L0225  Y   N
34A     QSFP DAC    FCI Electronics    1.0m 10131941-2010LF  2  CN1515FA102L0167  Y   N
35A     QSFP DAC    FCI Electronics    1.0m 10131941-2010LF  2  CN1515FA102L0101  Y   N
36A     QSFP DAC    FCI Electronics    1.0m 10131941-2010LF  2  CN1515FA102L0010  Y   N
37A     QSFP DAC    FCI Electronics    1.0m 10131941-2010LF  2  CN1449FA102L0029  Y   N
38A     QSFP DAC    FCI Electronics    2.0m 10142057-2020LF  E  CN2504ZE202L58B0  Y   N
39A     QSFP DAC    FCI Electronics    1.0m 10121178-2010LF  G  CN1539QV102L0022  Y   N
40A     QSFP DAC    FCI Electronics    1.0m 10131941-2010LF  2  CN1515FA102L0165  Y   N
41A     QSFP DAC    FCI Electronics    2.0m 10142057-2015LF  E  CN2440ZE152L4GW1  Y   N
42A     QSFP DAC    FCI Electronics    1.0m 10131941-2010LF  C  CN1638FA102L0098  Y   N
43A     QSFP DAC    FCI Electronics    3.0m 10142057-4030HLF E  CN2434ZE304H40M1  Y   N
44A     QSFP AOC    Hisense            7.0m DQF8503-4C07     02 S5189H7003E       Y   N    27.0
46A     QSFP AOC    Hisense            7.0m DQF8503-4C07     02 S5189H7003B       Y   N    34.0
47A     QSFP DAC    Amphenol           1.0m NJAAKR-CN01      A  APF2446N0172C6    Y   Y
48A     QSFP AOC    Hisense            7.0m DQF8503-4C07     02 S5189H7004F       Y   N    30.6
```

To correct, use `port config modify speed` to change the speed to a cable-supported value.

Refer to the *CN5000 Commands Guide* for “hardware” and “port”.

4.1.5. Errors During Switch Boot

4.1.5.1. Failed to Start Fan Sensor in Liquid-Cooled Switches

The firmware for CN5000 Switches is the same whether they are air-cooled or liquid-cooled. During start-up/boot of a liquid-cooled switch, you may see the following message:

```
[FAILED] Failed to start Fan Sensor.
```

Liquid-cooled switches do not have a Fan Sensor, therefore, you can ignore this message.

Note also that if you perform a `hardware check --verbose`, you may see an indication that the Fan Trays are "On", "Green", and "Good". These indications can be ignored as well.

4.2. Link Troubleshooting

4.2.1. Debugging Physical Link Issues

After you have run the proper commands and identified issues with links, it is useful to start root-causing the issues. This section focuses on the CN5000 Omni-Path Fabric physical links and not PCIe bus link issues.

The Omni-Path reporting tools are robust, but it can be confusing for new users to understand the difference between error counters and actual failures.

From an installation perspective, it is important to watch for physical issues with cabling, both copper and optical. In general, bend radius, cable insertion issues, and physical compression or damage to cables can result in transmission issues. The Omni-Path recovers from many issues transparently.

The following information can help you root-cause solid failures as well as marginal links. Most often the issue is resolved simply by re-installing a cable and verifying that it clicks into the connector socket on the SuperNIC or switch.

- View the QSFP/cable details of a specific switch port using the command:

```
opasmaquery -o cableinfo -d 10 -l <lid> -m <switch portnumber>
```

- To debug a particular switch, a useful technique is to get a snapshot of it using the command:

```
opareport -o snapshot -F portguid:value
```

Link issues may be the result of bad cables. Refer to the following sections for more information.

- Section 4.1.2 "Identifying Suspect Cables"
- Section 4.1.3 "Diagnosing Bad Cables"

4.2.1.1. Omni-Path Link Transition Flow

To debug link issues, it is helpful to understand the four key link states, starting from Offline and running properly in the final Active state.

NOTE

The Fabric Manager, `opafm`, must be running to transition physical links from the Init state to the Active state. If you subsequently stop the Fabric Manager when a link is in the Active state, the link remains active. You can safely make changes to the `opafm.xml` file for the Fabric Manager and restart the service without dropping active links.

PortState:

- Offline: Link down. QSFP not present or not visible to the SuperNIC driver.
- Polling: Physical link training in progress. At this point you do not know whether the other end of the QSFP is connected to a working Omni-Path device.
- Init: Link training has completed, both sides are present. Typically waiting for the Fabric Manager to enable the link.
- Active: Normal operating state of a fully functional link.

4.2.1.2. Verify the Fabric Manager is Running

From the Management Node, run the following command to report all SuperNICs and Switches.

```
# opafabricinfo
```

If it fails, try the following steps:

- Check the status of the Fabric Manager process using the command:

  ```
  # systemctl status opafm
  ```

- Restart the Fabric Manager using the command:

  ```
  # systemctl start opafm
  ```

4.2.1.3. Check the State of SuperNIC Links from a Server

If you are debugging server link issues, the `opainfo` command may be useful for a single server view.

`opainfo` captures a variety of data useful for debugging server-related link issues. Multiple Omni-Path commands can be used to extract individual data elements, however, this command is unique in the combination of data it provides.

- PortState: See Section 4.2.1.1 “Omni-Path Link Transition Flow”.
- LinkWidth: A fully functional link should indicate Act:4 and En:4.
- QSFP: Physical cable information for the QSFP, in this case a 5M Optical (AOC) Finisar cable.
- Link Quality: Range = 0 - 5 where 5 is Excellent.

```
# opainfo
hfi1_0:1                PortGID:0xfe80000000000000:001175010165b19c
   PortState:     Active
   LinkSpeed      Act: 25Gb         En: 25Gb
   LinkWidth      Act: 4            En: 4
```

```
   LinkWidthDnGrd ActTx: 4  Rx: 4   En: 3,4
   LCRC           Act: 14-bit       En: 14-bit,16-bit,48-bit Mgmt: True
   LID: 0x00000001-0x00000001       SM LID: 0x00000002 SL: 0
   QSFP: PassiveCu, 1m   FCI Electronics   P/N 10131941-2010LF  Rev 5
   Xmit Data:           22581581 MB Pkts:           5100825193
   Recv Data:           18725619 MB Pkts:           4024569756
   Link Quality: 5 (Excellent)
```

4.2.1.4. Link Width, Downgrades, and opafm.xml

By default, Omni-Path links run in x4 link width mode. Omni-Path has a highly robust link mechanism, as compared to InfiniBand, and it allows links to run in reduced widths with no data loss.

Three things to know:

1. By default, the `opafm.xml` configuration file requires links to start up in x4 link width mode. This is configurable separately for SuperNIC and ISL links using the **WidthPolicy** parameter.
2. Link downgrade ranges are also configurable in the `opafm.xml` file, using the **MaxDroppedLanes** parameter.
3. Default configuration example - a link that successfully starts up in x4 width and subsequently downgrades to x3 width continues to operate. If the link is restarted, by a server reboot, for example, and attempts to run by less than x4 width, then the link is disabled by the Fabric Manager and does not enter the Active state.

The `opainfo` command for SuperNICs is useful for checking the link width and link downgrade configuration on servers.

For a system view of all links that are running in less than x4 width mode, use the command:

```
# opareport -o errors -o slowlinks
```

4.2.1.5. How to Check Fabric Connectivity

For large fabrics, check the following flow in the topology spreadsheet:

- All host nodes should be defined as Type = FI in column F of the spreadsheet. All Edge switches (CN5000 Switches located on the edge of the network) should be defined as Type = SW in column L (destination from host to Edge) and column F (source for Edge to core that is also Edge switch). The following example shows links between host and Edge switch.

  ```
  R19 opahost1  1 FI  R19 opaedge1  13 SW opahost1_opae1p13 1m Cable CU
  ```

- All links between Edge switch to core that is also an Edge switch should be defined as Type = SW, as shown in the following example:

```
row1 rack01 opaedge1  1 SW row1 rack04 opaedgecore1 2 SW opae1p1_opac1p2 5M Cable Fiber
```

- All Director switches should be defined as Type = CL in column L (destination from Edge switch to Director switch). Column J (Name-2) should have the destination leaf and column K should have the port number on that leaf. The following example shows a link between an Edge switch to core that is a Director switch.

```
R19 opaedge1  5 SW  R72 opadirector1 01 L105B 11 CL opae1p5opad1L105Bp11 30m Fiber
```

- All Director Class Switches should be defined as shown in the following example:

```
Core Name:opadirector1 Core Group:row1 Core Rack:rack72 Core Size:1152 Core Full:0
```

- Set Core Full to 0 if the Director switch is not fully populated with all the leafs and spines. If it is fully populated, set Core Full to 1.

4.2.1.6. Link Debug CLI Commands

Task	CLI Command
Identify fabric errors.	`opareport -o errors`
Identify slow links (< x4 width).	`opareport -o slowlinks`
Obtain the LID of the switch.	Use `opaextractlids` for the links.
If a link is not coming up as Active, first bounce the link, then check the link state.	`opaportconfig -l <lid> -m <port> bounce`
Get detailed link info for all nodes connected to a Switch (edge) or leaf and their neighbor.	`opareport -M -m -A -s –o comps -d 20 -F lid:<lid of edge switch>:node`
Find links that are not plugged in or not seen by the interface. Find all links stuck in the Offline state.	`opareport -A -m -F portphysstate:offline -o comps -d 5`
Find all links stuck in the Polling state. **NOTE:** A link stuck in Polling may indicate that the other end of the cable is not inserted correctly. In this case, typically, one end is Polling and the other end is Offline.	`opareport -A -m -F portphysstate:polling -o comps -d 5`
To bounce a link, simulate a cable pull and re-insert on a server. **NOTE:** It may take up to 60 seconds for the port to re-enter the active state.	`opaportconfig bounce`
Check status of local SuperNIC ports.	`opainfo`
`opaportconfig` and `opaportinfo` are key commands for port debugging.	Run the commands with the `-help` option to see available parameters.

4.2.2. Link Down Reason

In order for two link partners to communicate reliably, a set of link states is defined to identify the ability of the link to move management or data traffic. Two indicators provide information about the reason a link went down:

- ***LinkDownReason*:** The reason the local port initiated a *LinkDown* from either the *LinkInit*, *LinkArmed*, or *LinkActive* state. It only captures the first reason for why the link is down (if more than one reason before the indicator is cleared).
- ***NeighborLinkDownReason*:** The value received from the neighbor.

NOTE

The SM is in charge of clearing both values to permit subsequent reasons to be recorded. The SM clears both these values as part of bringing the link to Armed.

The `opasaquery` tool can be used to show both the *LinkDownReason* and the *NeighborLinkDownReason*: `opasaquery -o portinfo`.

You can also look at the *LinkDownErrorLog*, which stores the last eight historical reasons for why the port went down, using: `opasaquery -o portinfo -vvv`.

Other tools showing the *LinkDownReason* are `opaportinfo` and `opasmaquery`.

The table below describes the *LinkDownReason* values:

Value	Description
0: None	
Corresponding to locally initiated link bounce due to PortErrorAction	
2: Bad Packet Length	Illegal packet length in the header
3: Packet Too Long	Packet longer than length
4: Packet Too Short	Packet shorter than length with normal tail
5: Bad source LID	Illegal SLID (0, using multicast as SLID. Does not include security validation of SLID)
6: Bad destination LID	Illegal DLID (0, does not match SuperNIC, multicast DLID on SC15)
7: Bad L2	Illegal L2 opcode
8: Bad SC	Unconfigured SC
10: Bad Mid Tail	Body/Tail received without a corresponding Head flit
12: Preempt Error	Preempting with same VL
13: Preempt VL15	Preempting a VL15 packet
14: Bad VL Marker	
17: Bad Head Distance	Distance violation between two head flits
18: Bad Tail Distance	Distance violation between two tail flits

Value	Description
19: Bad Control Distance	Distance violation between two credit LF command flits
20: Bad Credit Ack	Credits return for unsupported VL
21: Unsupported VL Marker	
22: Bad Preempt	Exceeding the interleaving level
23: Bad Control Flit	Unknown or reserved control flit received—deprecated
24: Exceed Multicast Limit	
32: Excessive Buffer Overrun	
Corresponding to local initiated intentional link down	
33: Unknown	
35: Reboot	Reboot or service reset
36: Neighbor Unknown	Link down was not locally initiated but no *LinkGoingDown* idle flit was received
39: FM Bounce	FM initiated bounce by transitioning from *LinkUp* to *Polling*.
40: Speed Policy	Link outside link policy
41: Width Policy	Link downgrade outside policy
Corresponding to local initiated intentional link down via transition to Offline or Disabled	
49: Disconnected	Link can never reach *LinkUp*
50: No Local Media Installed	Module is not installed in local port connector
51: Not Installed	Internal link not installed, due to absence of link partner FRU or backplane
52: Chassis Config	Chassis management forcing port *Offline* due to incompatible or absent link partner FRU or backplane
54: End to End not Installed	Silicon photonics mid-board module installed, but unable to detect link partner silicon photonics, due to absence of some part of the optical interconnect or absence of the remote module
56: Power Policy	Unable to enable port without exceeding power policy
57: Link Speed Policy	Link Speed Enabled policy is not able to be met due to a persistent cause
58: Link Width Policy	Link Width Enabled policy is not able to be met due to a persistent cause such as board design having insufficient lanes. Does not include dynamic reasons such as failed link negotiation or *LinkWidthDowngrade* below policy
60: Switch Management	User disabled via switch management interface (CLI, SNMP, Config file, etc.)
61: SMA Disabled	User disabled via SMA packet changing Physics Port State to *Disabled*
63: Transient	Port recently entered *Offline* and is waiting for a Timeout to ensure synchronization with link partner Physics Port State machine

If the link is currently down, the *LinkDownReason* and *LinkDownErrorLog* will not be available. It will be populated when the link comes back up.

In many cases, the *LinkDownReason* will be the same as the neighbor ports value of *NeighborLinkDownReason* and vice versa. However, there are exceptions.

The following table shows the *LinkDownReasons* that are only applicable to *LinkDownReason* and will not be used for *NeighborLinkDownReason*.

Value	Description
36	Neighbor Unknown
49	Disconnected
50	No Local Media Installed
51	Not Installed
54	End to End not Installed

The following are some sample combinations of values for a configuration with Device A connected to Device B:

- For a link down initiated by device A and device A is able to send the reason:

  A.LinkDownReason=X, A.NeighborLinkDownReason=0; B.LinkDownReason=0, B.NeighborLinkDownReason=X

- For a link down initiated by device A and B concurrently, where one of the devices is able to send the reason to the other device:

  A.LinkDownReason=X, A.NeighborLinkDownReason=Y; B.LinkDownReason=Y, B.NeighborLinkDownReason=X

- For a link down initiated by device A and device A is unable to send reason to device B:

  A.LinkDownReason=X, A.NeighborLinkDownReason=0; B.LinkDownReason=36 (Neighbor Unknown), B.NeighborLinkDownReason=0

- For an unexplained link down and device A or B is unable to send a reason code for the link going down (for example, cable failure):

  A.LinkDownReason=36 (Neighbor Unknown), A.NeighborLinkDownReason=0; B.LinkDownReason=36 (Neighbor Unknown), B.NeighborLinkDownReason=0

- For a link down initiated by device A due to hard failure (for example, power loss, hard reset, ASIC fault, FW/driver crash, etc) and device A is unable to send reason to device B:

  A.reason codes are inaccessible, powers back up as 0,0; B.LinkDownReason=36 (Neighbor Unknown), B.NeighborLinkDownReason=0

4.2.3. Port Type Information

Ports contain data that describe their types and the properties of the cables connected to them. This includes physical information as well as metadata about the port/cable.

Ports have one of the following types:

Port Type	Description
Disconnected	Part of design but physically unused
QSFP	Standard cable port
Fixed	Hardwired unchangeable port
Variable	Hardwired changeable port

4.2.3.1. Cable Information

A description of cable info is shown in the table below. The **bold** entries are most useful for diagnostics, while the remaining entries may aid in debugging. The amount of information shown, that is, the number of fields, can be controlled by a command line option to the tools.

NOTE

Different output levels may show the information with slightly altered labels.

Cable Info Item	Description
Identifier	Identifier type of cable
Power Class	Power Consumption Class: • 0 --- Power Class 1 (1.5 W max) • 1 --- Power Class 2 (2.0 W max) • 2 --- Power Class 3 (2.5 W max) • 3 --- Power Class 4 (3.5 W max)
Connector	Connector type code
NominalBR	Nominal supported bit rate
OM4Length	Supported OM4 fiber length
DeviceTech	Cable type description
VendorName	Vendor name
VendorOUI	Vendor OUI
VendorPN	Vendor part number
VendorRev	Vendor revision
CC_BASE	Checksum
VendorSN	Vendor serial number
DateCode	Vendor manufacturing date code
CC_EXT	Checksum

4.2.3.2. Port and Cable Verification Tools

In addition to `opareport`, several lower-level tools are useful for verifying port information and settings. These tools also allow you to view the properties for a cable in a link. These tools include:

- `opaportinfo`
- `opainfo`
- `opasaquery`
- `opasmaquery`

In some cases, the tools overlap with each other.

For additional details on each tool, refer to the *CN5000 Commands Guide*.

`opasaquery` allows you to query the Fabric Manager for its internal information. `opaportinfo`, `opainfo`, and `opasmaquery` can be used to query a port directly. `opainfo` provides information about local ports, whereas the other tools provide information about any port in the fabric.

For example:

- `opainfo -d4`: Provides basic port information along with detailed cable information for local port
- `opaportinfo -l2`: Provides port info for port with LID 2
- `opasmaquery -o portinfo`: Provides local port information
- `opasmaquery -o cableinfo -d4`: Provides local port cable information in detail
- `opasaquery -o portinfo`: Provides port information for all ports in fabric
- `opasaquery -o cableinfo`: Provides cable info

NOTE

Because `opareport` shows port names instead of LIDs, it may be more useful than using `opasaquery`.

4.2.4. Broken Intermediate Link

Sometimes message traffic passes through the fabric while other traffic appears to be blocked. In this case, MPI jobs fail to run.

In large cluster configurations, switches may be attached to other switches to supply the necessary inter-node connectivity. Problems with these inter-switch (or intermediate) links are sometimes more difficult to diagnose than failure of the final link between a switch and a node. The failure of an intermediate link may allow some traffic to pass through the fabric while other traffic is blocked or degraded.

If you notice this behavior in a multi-layer fabric, check that all switch cable connections are correct. Statistics for managed switches are available on a per-port basis, and may help with debugging. See your switch vendor for more information.

4.2.5. Port Counters

Each port in a CN5000 Omni-Path Fabric maintains a set of port counters to indicate both traffic and error counts. These counters can be grouped into the categories described in this section. Each port stops incrementing when the max value is reached, irrespective of counter size. Most of the counters are 64-bits in size. Exceptions are noted.

4.2.5.1. Utilization

These counters reflect the normal utilization of the port and Virtual Lane (VL) when present.

Several of these counters are used during the calculation of Congestion, SMA Congestion, and the Bubble Categories. The Utilization metrics provide a way of giving some of the other counter's context by comparing them to the amount of data or packets that were transmitted or received.

4.2.5.1.1. PortXmitData (TxD) and PortVLXmitData[n]

These counters indicate the total number of fabric packet flits transmitted. This does not include idle nor other LF command flits.

4.2.5.1.2. PortRcvData (RxD) and PortVLRcvData[n]

These counters indicate the total number of fabric packet flits received.

4.2.5.1.3. PortMulticastXmitPkts (MTxP)

This counter indicates the number of multicast and collective packets transmitted.

4.2.5.1.4. PortMulticastRcvPkts (MRxP)

This counter indicates the number of multicast and collective packets received.

4.2.5.2. Link Integrity

These counters reflect errors in the Physical (PHY) and Link Layers, as well as errors in firmware. In some cases, these errors are benign and can be ignored. However, in other cases, excessive link integrity errors can indicate a hardware problem such as a poor connection, marginal cable, incorrect length/model cable for signal rate, or damaged/broken hardware, such as bad connectors.

When a bad packet is detected, one of these counters is incremented and the Link Layer may either discard or replay the packet.

During the link training sequence, assorted errors may be observed. This is a normal part of the link training and clock synchronization process. Hence, errors observed as part of rebooting nodes or moving cables should not be considered a problem.

The category is calculated as a weighted sum of the counters in the group, with the exception of *ExcessiveBufferOverrunErrors*. The counters report on the receive side of the link. However, the counter can indicate a problem on either side of the link.

4.2.5.2.1. Link Quality Indicator (LQI)

This is a status indicator, similar to the signal strength bar display on a mobile phone, that enumerates link quality as a range of 0-5, with 5 being very good. Values in the lower part of the range may indicate hardware problems with components such as ports and cables that surface as signal integrity issues, leading to performance and other problems. The LQI gives you an instantaneous view of a link's quality on every hardware port.

Table 14. Link Quality Values and Description

Link Quality Value	Description
5	Working at or above preferred link quality, no action needed.
3	Working at the low end of acceptable link quality, recommend corrective action on the next maintenance window.
2	Working below acceptable link quality, recommend timely corrective action.
1	Working far below acceptable link quality, recommend immediate corrective action.
0	Link down

NOTE

Corrective action entails diagnosing the hardware (links/cables and ports/devices). For example: Are the cables bad or improperly placed? Is the SuperNIC/switch responsive? Does rebooting the device/server fix the issue?

4.2.5.2.2. LocalLinkIntegrityErrors (LLI) Counter

This counter indicates the number of retries initiated by a link transfer layer receiver.

The retry rate is represented by the Link Quality Indicator. A link that is meeting performance requirements has a Link Quality of 5, which corresponds to 1000 or fewer replays per second.

4.2.5.2.3. PortRcvErrors (RxE) Counter

This counter indicates the total number of packets containing an error that were received by the port, including Link Layer protocol violations and malformed packets. It indicates possible misconfiguration of a port, either by the Subnet Manager (SM) or by user intervention. It can also indicate hardware issues or extremely poor link signal integrity.

4.2.5.2.4. ExcessiveBufferOverrunErrors (EBO) Counter

This counter, associated with credit management, indicates an input buffer overrun. It indicates possible misconfiguration of a port, either by the SM or by user intervention. It can also indicate hardware issues or extremely poor link signal integrity.

4.2.5.2.5. LinkErrorRecovery (LER) Counter

This counter indicates the number of times the link has successfully completed the link error recovery process.

Link Quality Indicator is the primary indicator for link quality to use. This counter is factored into the value reported for Link Quality Indicator. This counter may be non-zero for a properly functioning link.

4.2.5.2.6. LinkDowned (LD) Counter

This counter indicates the total number of times the port has failed the link error recovery process and downed the link. These events can cause disruptions to fabric traffic.

4.2.5.2.7. UncorrectableErrors (Unc) Counter

This counter indicates the number of unrecoverable device errors. This may indicate a defect in the reporting device.

4.2.5.2.8. FMConfigErrors (FMC) Counter

This counter reports inconsistent configurations of the low-level SMA on either side of the link. It indicates possible misconfiguration of a port, either by the SM or by user intervention.

4.2.5.3. Congestion

These counters reflect possible errors that indicate traffic congestion in the fabric.

When congestion or a packet that has seen congestion is detected, one of these counters is incremented, and then depending on the issue reported, the packet must wait. In an extreme case, the packet may time out and be dropped.

The category is calculated as a weighted sum of the counters in the context of the utilization counters. With the exception of PortRcvFECN, the counters are all reported on the transmit side of the link. In addition, PortRcvBECN is only taken if the local node is a SuperNIC. However, the counter could indicate a problem on either side of the link.

4.2.5.3.1. CongDiscards (CD) Counter

NOTE

Formerly known as "SwPortCongestion".

This switch-only counter indicates the number of packets that were discarded as unable to transmit due to timeouts.

4.2.5.3.2. PortRcvFECN (RxF) Counter

When a device receives a packet with the Forward Explicit Congestion Notification (FECN) bit set to one, this counter is incremented.

4.2.5.3.3. PortRcvBECN (RxB) Counter

When a device receives a packet with the Backward Explicit Congestion Notification (BECN) bit set to one, this counter is incremented.

4.2.5.3.4. PortMarkFECN (MkF) Counter

This counter indicates the total number of packets that were marked Forward Explicit Congestion Notification (FECN) by the transmitter due to congestion.

4.2.5.3.5. PortXmitTimeCong (TxTC) Counter

This counter indicates the total number of *flit times* that the port was in a congested state for any data VL.

4.2.5.3.6. PortXmitWait (TxW) Counter

This counter indicates the amount of time (in *flit times*) any virtual lane had data but was unable to transmit due to no credits available.

4.2.5.4. SMA Congestion

These counters reflect congestion in the fabric specific to communication between the Subnet Manager and Subnet Manager Agents (SMA) using the management VL (VL 15).

The category is calculated exactly as the Congestion category using the same weights and the correct VL15 utilization counters.

4.2.5.4.1. PortVLXmitWait[15] (VLTxW[15]) Counter

This counter behaves the same as *PortXmitWait,* but it is restricted to VL 15, which carries only SM traffic.

4.2.5.4.2. VLCongDiscards[15] (VLCD[15]) Counter

NOTE

Formerly known as "*SwPortVLCongestio*.

This counter behaves the same as *Cong Discards*, but it is restricted to VL 15, which carries only SM traffic.

4.2.5.4.3. PortVLRcvFECN[15] (VLRxF[15]) Counter

This counter behaves the same as *PortRcvFECN*, but it is restricted to VL 15, which carries only SM traffic.

4.2.5.4.4. PortVLRcvBECN[15] (VLRxB[15]) Counter

This counter behaves the same as PortRcvBECN, but it is restricted to VL 15, which carries only SM traffic.

4.2.5.4.5. PortVLXmitTimeCong[15] (VLTxTC[15]) Counter

This counter behaves the same as PortXmitTimeCong, but it is restricted to VL 15, which carries only SM traffic.

4.2.5.4.6. PortVLMarkFECN[15] (VLMkF[15]) Counter

This counter behaves the same as PortMarkFECN, but it is restricted to VL 15, which carries only SM traffic.

4.2.5.5. Bubble

These counters occur when an unexpected idle flit is transmitted or received.

The transmit port sends idle flits until it can continue sending the rest of the packet. The category is calculated as follows:

1. The maximum value between the sum of the XmitWastedBW and XmitWaitData or the neighbor's PortRcvBubble.
2. Then divide the previous value by the port's utilization to provide context.

4.2.5.5.1. PortXmitWastedBW (WBW) Counter

This counter indicates the number of *flit times* where one or more packets have been started but the transmitters are forced to send idles due to bubbles in the ingress stream. Also, the VLs that have data to be sent are not permitted to preempt the currently transmitting VL.

4.2.5.5.2. PortXmitWaitData (TxWD) Counter

This counter indicates the number of *flit times* where one or more packets have been started but interrupted due to bubbles in the ingress stream.

4.2.5.5.3. PortRcvBubble (RxBb) Counter

This counter indicates the total number of *flit times* where one or more packets have started to be received, but the receiver received idle flits from the wire.

4.2.5.6. Security

These counters reflect possible security problems in the fabric.

Security problems can occur if a PKey or SLID violation occurs at the port during the ingress or egress of a packet.

The category is calculated as the sum of the neighbor's PortRcvConstraintErrors and the local port's PortXmitConstraintErrors.

4.2.5.6.1. PortRcvConstraintErrors (RxCE)

This counter is incremented when partition key or source LID violations are detected in a received packet, indicating a possible security issue or misconfiguration of device security settings.

4.2.5.6.2. PortXmitConstraintErrors (TxCE)

This counter is incremented when partition key violations are detected in a packet attempting to be transmitted, indicating a possible security issue or misconfiguration of device security settings.

4.2.5.7. Routing

These counters reflect possible routing issues. When a routing issue occurs, the offending packet is dropped.

A typical cause of this error is the routing to a wrong egress port or an improper Service Channel (SC) mapping. These errors can be a side effect of a port or device going down while traffic was still in flight to or through the given port or device.

4.2.5.7.1. PortRcvSwitchRelayErrors (RxSR)

This counter indicates the number of packets that were dropped due to internal routing errors. It indicates possible misconfiguration of a switch by the SM.

4.2.5.8. Other

These counters do not fit into any of the previous categories.

4.2.5.8.1. PortRcvRemotePhysicalErrors (RxRP)

This counter indicates the number of downstream effects of signal integrity (SI) problems. It indicates an SI issue in the upstream path.

This counter is not included in the Link Integrity category calculation because it does not directly indicate the link that had the issue, so it can be misleading.

4.2.5.8.2. PortXmitDiscards (TxDc)

This counter indicates the number of packets dropped due to several reasons including timeouts and improper packet lengths.

NOTE

This counter is a super set that includes Congestion Discards counter.

4.3. Software Troubleshooting

4.3.1. Kernel and Initialization Issues

This section describes Issues that may prevent the system from coming up properly.

4.3.1.1. Driver Load Fails Due to Unsupported Kernel

If you try to load the Omni-Path driver on a kernel that the CN5000 OPX Software does not support, the load fails with error messages that point to `hfi1.ko`.

To correct this problem, install one of the appropriate supported Linux kernel versions, then reload the driver.

4.3.1.2. Rebuild or Reinstall Drivers if Different Kernel Installed

If you upgrade the kernel, you must reboot and then rebuild or reinstall the Omni-Path kernel modules (drivers). Refer to the *CN5000 Fabric Installation Guide* for more information.

4.3.1.3. Omni-Path Interrupts Not Working

The driver cannot configure the Omni-Path link to a usable state unless interrupts are working. Check for this problem with the command:

```
$ grep hfi1 /proc/interrupts
```

NOTE

The output you see may vary depending on board type, distribution, or update level, and the number of CPUs in the system.

If there is no output at all, the driver initialization failed. For more information on driver problems, see Section 4.3.1.1 “Driver Load Fails Due to Unsupported Kernel” or Section 4.3.1.5 “CN5000 Omni-Path SuperNIC Initialization Failure”.

If the output is similar to one of these lines, then interrupts are not being delivered to the driver.

```
-MSI-edge    hfi1_0 sdma6
177:    0    0   0    PCI-MSI-edge    hfi1_0 sdma7
178:    0    0   0    PCI-MSI-edge    hfi1_0 sdma8
179:    0    0   0    PCI-MSI-edge    hfi1_0 sdma9
180:    0    0   0    PCI-MSI-edge    hfi1_0 sdma10
181:    0    0   0    PCI-MSI-edge    hfi1_0 sdma11
182:    0    0   0    PCI-MSI-edge    hfi1_0 sdma12
183:    0    0   0    PCI-MSI-edge    hfi1_0 sdma13
184:    0    0   0    PCI-MSI-edge    hfi1_0 sdma14
185:    0    0   0    PCI-MSI-edge    hfi1_0 sdma15
```

```
186:   39    0   0    PCI-MSI-edge    hfi1_0 kctxt0
187:    1   77   0    PCI-MSI-edge    hfi1_0 kctxt1
188:    0    0   0    PCI-MSI-edge    hfi1_0 kctxt2
```

A zero count in all CPU columns means that no Omni-Path interrupts have been delivered to the processor.

The possible causes of this problem are:

- Booting the Linux kernel with ACPI disabled on either the boot command line or in the BIOS configuration.
- Other Omni-Path initialization failures.

To check if the kernel was booted with the `noacpi` or `pci=noacpi` option, use this command:

```
$ grep -i acpi /proc/cmdline
```

If output is displayed, fix the kernel boot command line so that ACPI is enabled. This command line can be set in various ways, depending on your OS distribution. If no output is displayed, check that ACPI is enabled in your BIOS settings.

To track down other initialization failures, see Section 4.3.1.5 "CN5000 Omni-Path SuperNIC Initialization Failure".

4.3.1.4. OpenFabrics Load Errors if SuperNIC Driver Load Fails

When the SuperNIC driver fails to load, the other OpenFabrics drivers/modules are loaded and shown by `lsmod`. However, commands such as `ibv_devinfo` fail if the SuperNIC driver fails to load, as shown in the following example:

```
ibv_devinfo
libibverbs: Fatal: couldn't read uverbs ABI version.
No Omni-Path devices found
```

4.3.1.5. CN5000 Omni-Path SuperNIC Initialization Failure

There may be cases where the SuperNIC driver was not properly initialized. Symptoms of this may show up in error messages from an MPI job or another program.

Here is a sample command and error message:

```
$ mpirun -np 2 -m ~/tmp/mbu13 osu_latency
<nodename>:hfi_userinit: assign_port command failed: Network is down
<nodename>:can't open /dev/hfi1, network down
```

This is followed by messages of this type after 60 seconds:

```
MPIRUN<node_where_started>: 1 rank has not yet exited 60 seconds after rank 0 (node
<nodename>) exited without reaching MPI_Finalize().
```

```
MPIRUN<node_where_started>:Waiting at most another 60 seconds for the remaining
ranks to do a clean shutdown before terminating 1 node processes.
```

If this error appears, check to see if the CN5000 Omni-Path SuperNIC driver is loaded with the command:

```
$ lsmod | grep hfi
```

If no output is displayed, the driver did not load for some reason. In this case, try the following commands (as root):

```
modprobe -v hfi1
lsmod | grep hfi1
dmesg | grep -i hfi1 | tail -25
```

The output indicates whether the driver has loaded or not. Printing out messages using `dmesg` may help to locate any problems with the SuperNIC driver.

If the driver loaded, but MPI or other programs are not working, check to see if problems were detected during the driver and hardware initialization with the command:

```
$ dmesg | grep -i hfi1
```

This command may generate more than one screen of output.

Also, check the link status with the command:

```
$ hfi1_control -iv
```

4.3.2. OpenFabrics and Omni-Path Issues

This section describes issues related to OpenFabrics, including Subnet Managers and Omni-Path.

4.3.2.1. Stop Services Before Stopping/Restarting CN5000 OPX Software

The Fabric Manager must be stopped before stopping, starting, or restarting the CN5000 OPX Software.

Use the `systemctl` command to stop or start the Fabric Manager:

```
# systemctl [start|stop|restart] opafm
```

To verify the status of the Fabric Manager, run the following command:

```
# systemctl status opafm
```

4.3.3. Troubleshooting the Fabric Manager

The Fabric Manager provides log messages for the following:

- Events (NOTICE)
- Information (INFO)
- Warning (WARN)
- Errors

4.3.3.1. Fabric Manager Event Messages

The Fabric Manager logs significant fabric events in a standard machine-readable format. The format for these special event messages provides information not only about the event, but information about what nodes in the fabric are causing the event.

The format of these messages is as follows:

```
<prefix>;MSG:<msgType>|SM:<sm_node_desc>:port <sm_port_number>|
COND:<condition>|NODE:<node_desc>:port <port_number>:<node_guid>|
LINKEDTO:<linked_desc>:port <linked_port>:<linked_guid>| DETAIL:<details>
```

Where:

- `<prefix>` – Includes the date and time information of the event along with either the slot number OR hostname and IP address of the Fabric Manager reporting the message.
- `<msgType>` – Is one of the following values:
  - `ERROR`
  - `WARNING`
  - `NOTICE`
  - `INFORMATION`
- `<sm_node_desc>` and `<sm_port_number>` – Indicate the node name and port number of the SM that is reporting the message, prefixed with the word '`port`'. Any pipes (`|`) or colons (`:`) in the node description will be converted to spaces in the log message.
- `<condition>` – Is one of the conditions from the event SM Reporting Table that is detailed in the Section 4.3.3.1.1 "Event Descriptions". The condition text includes a unique identification number.
- `<node_desc>`, `<port_number>`, and `<node_guid>` are the node description, port number, and node GUID of the port and node that are primarily responsible for the event. Any pipes (`|`) or colons (`:`) in the node description will be converted to spaces in the log message.
- `<linked_desc>`, `<linked_port>`, and `<linked_guid>` are optional fields describing the other end of the link. These fields and the '`LINKEDTO`' keyword will only be shown

in applicable messages. Any pipes (|) or colons (:) in the node description will be converted to spaces in the log message.

- <details> is an optional free-form field detailing additional information useful in diagnosing the log message cause.

4.3.3.1.1. Event Descriptions

The following sections describe the Fabric Manager event messages, their severity, an explanation, and possible causes for the event.

#1 Redundancy Lost

The subnet manager emits this message when it is the only running Subnet Manager on a given subnet.

- **Severity**

  Warning

- **Causes**

  No redundant SM exists on the subnet.

  A user shut down a redundant SM or possibly disconnected, or shut down, the node on which the SM was running.

- **Action**

  If running redundant SMs on a fabric, verify the health of each host or switch running an SM.

#2 Redundancy Restored

The Master SM for the subnet detected that another SM has come online.

- **Severity**

  Notice

- **Causes**

  A user started a redundant SM on another host or switch.

  A user just connected two separate subnets together.

- **Action**

  None

#3 Appearance in Fabric

A new SuperNIC port, switch, inter-switch link, or Subnet Manager was detected by the master Subnet Manager.

- **Severity**

  Notice

- **Causes**

  User action

- **Action**

  None

#4 Disappearance from Fabric

A SuperNIC port, switch, inter-switch link, or Subnet Manager has disappeared from fabric. This encompasses system shutdowns and loss of connectivity.

- **Severity**

  Notice

- **Action**

  The administrator should validate whether or not the components have disappeared from the fabric due to user action or not. Nodes will typically disappear from the fabric when they are rebooted, re-cabled, or if their Omni-Path Fabric stacks are stopped.

#5 SM State Change to Master

Subnet Manager transitioned into the 'master' state from one of the 'standby', 'discovering', or 'not active' states.

- **Severity**

  Notice

- **Action**

  The administrator should check the state of the machine (or chassis) that was providing the master SM service to determine if it has failed and needs to be replaced, or whether the state change occurred due to user action.

- **Example**

```
Nov 28 17:45:25 sample-host fm0_sm[29326]:
;MSG:NOTICE|SM:sample-host.sample-domain.com:port 1|COND:#5 SM state to
master|NODE:sample-host.sample-domain.com:port
1:0x0x00066a00a0000405|DETAIL:transition from DISCOVERING to MASTER
```

#6 SM State Change to Standby

Subnet Manager transitioned from 'master' into 'standby' state.

- **Severity**

  Notice

- **Action**

  The administrator should validate that this was due to a modification in the CN5000 Omni-Path Fabric network configuration. If not, then this issue should be reported to customer support.

- **Example**

```
Nov 29 12:15:28 sample-host fm0_sm[31247]:
;MSG:NOTICE|SM:sample-host.sample-domain.com:port 1|COND:#6 SM state to
standby|NODE:sample-host.sample-domain.com:port
1:0x0x00066a00a0000405|DETAIL:transition from MASTER to STANDBY
```

#7 SM Shutdown

The master Subnet Manager is shutting down.

- **Severity**

  Notice

- **Action**

  The administrator should check the state of the machine (or chassis) that was providing the master SM service, or whether the state change occurred due to user action.

- **Example**

```
;MSG:NOTICE|SM:sample-host.sample-domain.com:port 1|COND:#7 SM
shutdown|NODE:sample-host.sample-domain.com:port 1:0x0x00066a00a0000405|DETAIL:
```

#8 Fabric Initialization Error

Some form of error occurred during fabric initialization.

- **Severity**

  Notice

- **Explanation**

  Examples of possible errors include:

  – Link could not be activated in 4x mode.

  – Subnet Manager could not initialize a port or node with proper configuration.

- **Action**

  The administrator should perform the fabric troubleshooting procedure to isolate and repair the faulty component. The faulty component could be the SM platform itself (for example, its own SuperNIC) or a component in the CN5000 Omni-Path Fabric network.

- **Example**

```
Apr  6 22:48:42 sample-host fm0_sm[21458]: sample-host; MSG:NOTICE|SM:sample-host:port
2|COND:#8
Fabric initialization error|NODE:sample-host2:port
1:0x0011750000ffd7af|LINKEDTO:Cornelis OPA Switch:port
18:0x00066a00d9000108|DETAIL:Failed to set portinfo for node
```

#9 Link Integrity Error

The SM received an asynchronous trap from a switch or end-port indicating a link integrity problem.

- **Severity**

  Notice

- **Action**

  The administrator should perform the fabric troubleshooting procedure to isolate and repair the faulty component. This is typically due to a bad cable, an incorrect cable being used for the signaling rate and cable length (for example, too small a wire gauge), or a hardware failure on one of the two SuperNIC ports.

#10 Security Error

The SM received an asynchronous trap from a switch or end-port indicating a management key violation.

- **Severity**

  Notice

- **Action**

  The administrator should validate that the software configuration has not changed, because this issue is most likely due to a configuration issue. However, this event could also indicate a more serious issue such as a hacking attempt.

#11 Other Exception

The Subnet Manager encountered an error at some time after fabric initialization.

- **Severity**

  Notice

- **Explanation**

  Examples of possible errors are:

  – The SM received an invalid request for information.

  – The SM could not perform the action requested by another fabric entity such as a request to create or join a multicast group with an unrealizable MTU or rate.

- **Action**

  The administrator should check to see if other SM-related problems have occurred and perform the corrective actions for those items. If these other exceptions continue to persist, then customer support should be contacted.

#12 Fabric Summary

A brief message describing the number of changes that the SM detected on its last subnet sweep. This message will include totals for the number of switches, SuperNICs, end-ports, total physical ports, and SMs that have appeared or disappeared from the fabric. This message will only be logged at the end of a subnet sweep if the SM had detected changes.

- **Severity**

  Notice

- **Action**

  As this is only a summary of events detected during a fabric sweep, the administrator should examine the logs for preceding messages that describe the fabric changes in detail.

- **Example**

```
Apr  8 15:31:36 sample-host fm0_sm[21458]: sample-host; MSG:NOTICE|SM:sample-host:port
2|COND:#12 Fabric Summary|NODE:sample-host:port 2:0x00066a01a0000405|DETAIL:Change
Summary: 1 SWs disappeared, 0 HFIs appeared, 1 end ports disappeared, 3 total ports
disappeared, 0 SMs appeared
```

#13 SM State Change to Inactive

Subnet Manager transitioned from *standby* into *inactive* state.

- **Severity**

  Notice

- **Action**

  The administrator should check for inconsistencies in XML configurations between the master SM and this SM.

- **Example**

```
Nov 29 12:15:28 sample-host fm0_sm[31247]:
;MSG:NOTICE|SM:sample-host.sample-domain.com:port 1|COND:#13 SM state to
inactive|NODE:sample-host.sample-domain.com:port
1:0x0x00066a00a0000405|DETAIL:transition from STANDBY to NOTACTIVE
```

#14 SM Inconsistency

Deactivating the Standby Subnet Manager and Secondary Performance Manager due to inconsistent Subnet Manager XML configuration on Standby.

- **Severity**

  Warning

- **Action**

  If the condition persists, compare the XML configuration files between the master and standby SM for inconsistencies.

- **Example**

```
Oct 22 12:49:06 shaggy fm0_sm[31032]: shaggy; MSG:WARNING|SM:shaggy:port 1|COND:#14
SM standby configuration inconsistency|NODE:i9k118:port
0:0x00066a00d8000118|DETAIL:Deactivating standby SM i9k118 : 0x00066a00d8000118
which has a SM configuration inconsistency with master! The secondary PM will also
be deactivated.
```

#15 SM Virtual Fabric Inconsistency

Deactivating the Standby Subnet Manager and Secondary Performance Manager due to inconsistent Subnet Manager Virtual Fabrics XML configuration on Standby.

- **Severity**

  Warning

- **Action**

  If the condition persists, compare the XML configuration files between the master and standby SM for inconsistencies.

- **Example**

```
Oct 22 12:23:41 shaggy fm0_sm[30778]: shaggy; MSG:WARNING|SM:shaggy:port 1|COND:#15
SM standby virtual fabric configuration inconsistency|NODE:i9k118:port
0:0x00066a00d8000118|DETAIL:Deactivating standby SM i9k118 : 0x00066a00d8000118
which has a Virtual Fabric configuration inconsistency with master! The secondary
PM will also be deactivated.
```

#16 Reserved for Future Use

Reserved

#17 PM Inconsistency

Deactivating the Secondary Performance Manager and Standby Subnet Manager due to inconsistent Performance Manager XML configuration on Secondary.

- **Severity**

  Warning

- **Action**

  If the condition persists, compare the XML configuration files between the primary and secondary PM for inconsistencies.

- **Example**

```
Oct 22 12:51:42 shaggy fm0_sm[31173]: shaggy; MSG:WARNING|SM:shaggy:port 1|COND:#17
PM secondary configuration inconsistency|NODE:i9k118:port
0:0x00066a00d8000118|DETAIL:Attempting to deactivate secondary PM which has a
configuration inconsistency with primary! The standby SM will also be deactivated.
```

4.3.3.2. Other Log Messages

In addition to the Fabric Manager Event messages detailed in the previous section, the Fabric Manager software suite may emit other log messages that provide extra detail for use by technical personnel in troubleshooting fabric issues.

Log messages generally follow the following format:

```
<prefix>: <severity>[<module>]: <component>: <function>: <message>
```

Where:

- `prefix` - Includes time and date followed by the hostname and/or IP of the Fabric Manager reporting the message, the instance name, and a Process ID (PID) number.
- `severity` - One of the following:

  `FATAL`, `ERROR`, `WARN`, `NOTIC`, `INFO`, `PROGR`, `VBOSE`, `DBG[1-4]`, `ENTER`, or `EXIT`.

- `module` - Program module that generated the message. Typically the name of the sub-component or library that saw the event.
- `component` - Name of the Fabric Manager process that owns the module.
- `function` - Part of the sub-module where the event occurred. This is probably only useful for developers but might give insight into what the Fabric Manager is currently doing.
- `message` - Free form message text giving more details or explaining the event.

NOTE

Some of the listed components of the formatting may be omitted.

Example:

```
Jan 20 14:40:22 phgppriv36 fm0_sm[4082]: PROGR[topology]: SM:
topology_main: DISCOVERY CYCLE END. 0 SWs, 2 HFIs, 2 end ports,
2 total ports, 1 SM(s), 26 packets, 0 retries, 0.004 sec sweep
```

4.3.3.2.1. Information (INFO)

Last full member of multicast group GID 0xff12401bffff0000:00000000ffffffff is no longer in fabric, deleting all members

- **SM Area**

  Discovery
- **Meaning**

  The last full member of the group has left. The group is removed from the fabric.
- **Action**

  None

topology_discovery: now running as a STANDBY SM

- **SM Area**

  Discovery
- **Meaning**

  SM has transitioned to STANDBY mode.
- **Action**

  None

TT: DISCOVERY CYCLE START

- **SM Area**

  Discovery
- **Meaning**

  Discovery sweep has started.
- **Action**

  None

TT: DISCOVERY CYCLE END

- **SM Area**

  Discovery
- **Meaning**

  Discovery sweep has ended.
- **Action**

  None

Port x of node [y] HFI1 belongs to another SM [0x0001]; Marking port as NOT MINE!

- **SM Area**

  Discovery
- **Meaning**

  Usually happens during the merging of two fabrics.
- **Action**

  None

sa_PathRecord: requested source GUID/LID not found/active in current topology

- **SM Area**

  Administrator
- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric.
- **Action**

  Check the health of the requester and the connected port if the message persists.

sa_PathRecord: requested destination GUID not an active port nor a Multicast Group

- **SM Area**

  Administrator

- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric or the destination has dropped from fabric.

- **Action**

  None

sa_XXXXXXXX: Cannot find source lid of 0x0001 in topology in request to subscribe/unsubscribe...

- **SM Area**

  Administrator

- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric or a request received from the node that the SM has dropped from the fabric due to non-response to SMA queries.

- **Action**

  Check the health of the node at lid 0x0001 if the fabric is stable.

sa_XXXXXXXX: requested source Lid/GUID not found/active in current topology

- **SM Area**

  Administrator

- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric or request received from the node that the SM has dropped from the fabric due to non-response to SMA queries.

- **Action**

  Check the health of node at lid 0x0001 if the fabric is stable.

sa_McMemberRecord_Set: Port GID in request (0xFE80000000000000:00066a00d9000143) from HFI1, Port 0x00066a00d9000143, LID 0x0001, for group 0xFF12401BFFFF0000:00000000FFFFFFFF can't be found or not active in current topology, returning status 0x0001/0x0200

- **SM Area**

  Administrator

- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric or a request received from the node that the SM has dropped from the fabric due to the non-response to SMA queries.

- **Action**

  Check the health of node at lid 0x0001 if fabric is stable.

sa_McMemberRecord_Set: Last full member left multicast group GID 0xFF12401BFFFF0000:00000000FFFFFFFF, deleting group and all members

- **SM Area**

  Administrator

- **Meaning**

  Group is cleaned out when last the FULL member leaves.

- **Action**

  None

4.3.3.2.2. Warning (WARN)

failed to send reply [status=x] to SMInfo GET request from node HFI1 guid 0x00066a00d9000143, TID=0x811E796027000000

- **SM Area**

  SM to SM Communication

- **Meaning**

  Lost communication path to other SM on node HFI1.

- **Action**

  Check the health of the node described in the message and the status of the SM node.

failed to send reply [status=x] to SMInfo SET request from node HFI1 guid 0x00066a00d9000143, TID=0x811E796027000000

- **SM Area**

  SM to SM Communication

- **Meaning**

  Lost communication path to the other SM on node HFI1.

- **Action**

  Check the health of the node described in the message and the status of the SM node.

SmInfo SET control packet not from a Master SM on node HFI1, lid [0x1], guid 0x00066a00d9000143, TID=0x811E796027000000

- **SM Area**

  SM to SM Communication

- **Meaning**

  The SM on node HFI1 is violating the protocol.

- **Action**

  If the condition persists, turn off the SM on node HFI1.

Standby SM received invalid AMOD[1-5] from SM node HFI1, LID [0x1], guid [0x00066a00d9000143], TID=0x811E796027000000

- **SM Area**

  SM to SM Communication

- **Meaning**

  SM on node HFI1 is violating the protocol specification.

- **Action**

  If the condition persists, turn off the SM on node HFI1.

MASTER SM did not receive response to Handover Acknowledgement from SM node HFI1, LID [0x1], guid [0x00066a00d9000143]

- **SM Area**

  SM to SM Communication

- **Meaning**

  The SM on node HFI1 is incompatible or lost the communication path.

- **Action**

  Remove the incompatible SM from the fabric or check the health of the node HFI1.

INACTIVE SM received invalid STANDBY transition request from SM node HFI1, LID [0x1], guid [0x00066a00d9000143], TID=0x811E796027000000

- **SM Area**

  SM to SM Communication

- **Meaning**

  The SM on node HFI1 is violating the protocol specification.

- **Action**

  If the condition persists, turn off the SM on node HFI1.

Master SM received invalid Handover Ack from remote SM HFI1, LID [0x1], guid [0x00066a00d9000143], TID=0x811E796027000000; remote not in STANDBY state [Discovering]

- **SM Area**

  SM to SM Communication

- **Meaning**

  The SM on node HFI1 is violating the protocol specification.

- **Action**

  If the condition persists, turn off the SM on node HFI1.

Master SM received invalid MASTER transition [requested state] from remote [remote state] SM HFI1, LID [0x1], guid [0x00066a00d9000143], TID=0x811E796027000000

- **SM Area**

  SM to SM Communication

- **Meaning**

  The SM on node HFI1 is violating the protocol specification.

- **Action**

  If the condition persists, turn off the SM on node HFI1.

Master SM did not receive response to Handover Acknowledgment from [remote state] SM node HFI1, LID [0x1], guid [0x00066a00d9000143]

- **SM Area**

  SM to SM Communication

- **Meaning**

  Lost communication path to the other SM on node HFI1.

- **Action**

  Check the health of the node described in the message and the status of the SM node.

SM at shaggy HFI-1, portGuid=0x0011750000ff8f4d has a different SM configuration consistency checksum [418863] from us [417845]

- **SM Area**

  SM to SM Communication

- **Meaning**

  The SM on node HFI1 configuration does not match the master.

- **Action**

  Verify that the XML configuration between master and standby SM is consistent.

No transitions allowed from DISCOVERING state; Got (ANY) request from [state] SM node HFI1, LID [0x1], guid [0x00066a00d9000143]

- **SM Area**

  SM to SM Communication

- **Meaning**

  The SM on node HFI1 is violating the protocol specification.

- **Action**

  If the condition persists turn off the SM on node HFI1.

SmInfo from SM at SMLID[0x1] indicates SM is no longer master, switching to DISCOVERY state

- **SM Area**

  SM to SM Communication

- **Meaning**

  Remote SM may have handed over to another SM on the fabric.

- **Action**

  None

Switching to DISCOVERY state; Failed to get SmInfo from master SM at LID 0x1

- **SM Area**

  SM to SM Communication

- **Meaning**

  Lost the communication path to the other SM at lid 0x1.

- **Action**

  Check the health of the node described in the message and the status of the SM node.

too many errors during sweep, will re-sweep in a few seconds

- **SM Area**

  Discovery

- **Meaning**

  Multiple cable pulls or chassis removal/insertion event.

- **Action**

  Check links with high error count and reseat or replace cable. If the condition persists, capture the log information and call support.

unable to setup port [x] of node Sw1/HFI1, nodeGuid 0x00066a00d9000143, ignoring port!

- **SM Area**

  Discovery

- **Meaning**

  Lost communication path to node HFI1.

- **Action**

  Check the health of the node port described in the message.

Get NodeInfo failed for node off Port x of Node 0x00066a00d9000143:HFI1, status=7

- **SM Area**

  Discovery

- **Meaning**

  The node connected to port x of HFI1 is not responding.

- **Action**

  Check the health of the node connected to the port.

Get NodeDesc failed for node off Port X of Node 0x00066a00d9000143:HFI1, status = 7

- **SM Area**

  Discovery

- **Meaning**

  The node connected to port x of HFI1 is not responding.

- **Action**

  Check the health of the node connected to the port.

Failed to get Switchinfo for node sw1 guid 0x00066a00d9000143: status = 7

- **SM Area**

  Discovery

- **Meaning**

  Switch node 1 is not responding.

- **Action**

  If the condition persists, check the health of the switch and capture health data if possible.

Failed to set Switchinfo for node sw1 nodeGuid 0x00066a00d9000143: status = 7

- **SM Area**

  Discovery

- **Meaning**

  Switch node 1 not responding.

- **Action**

  If the condition persists, check the health of the switch and capture health data if possible.

Failed to get PortInfo from NodeGUID x [Hfi1] Port 1; Ignoring port!

- **SM Area**

  Discovery

- **Meaning**

  Port x of HFI1 not responding.

- **Action**

  Check the health of node HFI1.

port on other side of node sw1 index x port X is not active

- **SM Area**

  Discovery

- **Meaning**

  The node may have been marked down if it did not respond to SMA queries.

- **Action**

  Check the health of the node connected to switch 1 port X.

Node Hfi1 [0x00066a00d9000143] port[x] returned MKEY[0x1] when MKEY[0x0] was requested!

- **SM Area**

  Discovery

- **Meaning**

  Another SM with a different Mkey configured.

- **Action**

  Stop one of the subnet managers and make the configuration consistent.

Cannot get PORTINFO for node HFI1 nodeGuid 0x00066a00d9000143 port X status=Y

- **SM Area**

  Discovery

- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric. Persistence indicates that the node may be having problems.

- **Action**

  Check the health of the node HFI1/SW1 if the fabric was idle or persistent condition.

Cannot set PORTINFO for node HFI1 nodeGuid 0x00066a00d9000143 port X status=Y

- **SM Area**

  Discovery

- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric. Persistence indicates the node may be having problems.

- **Action**

  Check the health of node HFI1/SW1 if the fabric was idle or persistent condition.

Could not find neighbor for NodeGUID x Sw1] (neighbor idx x, port y) in new topology; spanning tree not up to date

- **SM Area**

  Discovery

- **Meaning**

  Caused by simultaneous removal/insertion events in the fabric.

- **Action**

  None

sa_NodeRecord_GetTable: Invalid node type[~1-3] in request from lid 0x1

- **SM Area**

  Administrator

- **Meaning**

  Invalid data in the SA request.

- **Action**

  Check the health of the requester at lid 0x1.

sa_PathRecord_Set: Cannot find path to port 0x00066a00d9000144 from port 0x00066a00d9000143: failing src/dst pkey 0x800d validation

- **SM Area**

  Administrator

- **Meaning**

  The source and destination do not share a partition with the given PKey.

- **Action**

  Configuration change required if they should have access.

sa_PathRecord_Set: Cannot find path to port 0x00066a00d9000144 from port 0x00066a00d9000143: failing req/dst pkey validation

- **SM Area**

  Administrator

- **Meaning**

  The requesting node and destination do not share a PKey.

- **Action**

  Configuration change required if they should have access.

sa_PathRecord_Set: Cannot find path to port 0x00066a00d9000144 from port 0x00066a00d9000143: failing vFabric rate validation (mtu=2,rate=3)

- **SM Area**

  Administrator

- **Meaning**

  The source and destination do not share a path in a vFabric that contains limitations on MTU and rate.

- **Action**

  Configuration change required if the path is valid.

sa_PathRecord/SA_TraceRecord: Failed PKey check for source x and destination y for PKey 0x800d

- **SM Area**

  Administrator

- **Meaning**

  A request was for a given Pkey, but the source of the query is not a member of the same partition.

- **Action**

  Configuration change may be necessary.

sa_PathRecord/SA_TraceRecord: Failed pairwise PKey check for request

- **SM Area**

  Administrator

- **Meaning**

  A query request failed pairwise PKey checks.

- **Action**

  Configuration change may be necessary.

sm_resolve_pkeys_for_vfs: VFabric has undefined pkey. Assigning pkey 0x3

- **SM Area**

  Configuration
- **Meaning**

  A vFabric with an undefined PKey has been assigned a PKey.
- **Action**

  None

sa_ServiceRecord_GetTable: Filter serviced record ID=0x1000000000003531 from lid 0x4 due to pkey mismatch from request port

- **SM Area**

  Configuration
- **Meaning**

  PKey validation failed for service record, request node does not have valid PKey.
- **Action**

  Configuration change required if request should be valid.

sa_xxxxxx: too many records for SA_CM_GET

- **SM Area**

  Administrator
- **Meaning**

  May have duplicate data in the fabric.
- **Action**

  Check topology data for duplicate GUIDs.

sa_TraceRecord/Pathrecord_set: Cannot find path to port 0x00066a00d9000144 from port 0x00066a00d9000144: LFT entry for destination is 255 from switch Sw1 (nodeGuid 0x00066a00d9000999)

- **SM Area**

  Administrator

- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric.

- **Action**

  Check SW1 for a bad port and the health of the destination node if the condition persists.

sa_TraceRecord/Pathrecord_set: Cannot find path to destination port 0x00066a00d9000144 from source port 0x00066a00d9000143; INVALID TOPOLOGY, next/last_nodep is NULL

- **SM Area**

  Administrator

- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric.

- **Action**

  Check SW1 for a bad port and health of the destination node if the condition persists.

sa_updateMcDeleteCountForPort: MC Dos threshold exceeded for: Node= HFI1, GUID=0x00066a00d9000143, PortIndex=1; bouncing port

- **SM Area**

  Administrator

- **Meaning**

  SM Multicast denial of service configured and the threshold has been reached. Bouncing the port in an attempt to clear the issue.

- **Action**

  If multiple occurrences, check the health of the node HFI1.

sa_updateMcDeleteCountForPort: MC Dos threshold exceeded for: Node= HFI1, GUID=0x00066a00d9000143, PortIndex=1; disabling port

- **SM Area**

  Administrator

- **Meaning**

  SM Multicast denial of service configured and the threshold has been reached, disabling the port.

- **Action**

  Check the health of the node HFI1.

4.3.3.2.3. ERROR

could not perform HANDOVER to remote SM HFI1: 0x00066a00d9000143

- **SM Area**

  SM to SM communication

- **Meaning**

  Lost communication path to the other SM on node GUID 0x00066a00d9000143.

- **Action**

  Check the health of the node HFI1 and the status of the SM node.

can't get PortInfo, sleeping

- **SM Area**

  Discovery

- **Meaning**

  topology_initialize: cannot get PortInfo; sleeping.

- **Action**

  Make sure the stack is running. Restart the SM node and stack.

port state < INIT, sleeping

- **SM Area**

  Discovery
- **Meaning**

  The node port of the SM is down.
- **Action**

  Be certain the host cable is connected to a switch (host SM only).

can't get/set isSM, sleeping

- **SM Area**

  Discovery
- **Meaning**

  SM cannot communicate with the stack.
- **Action**

  Make sure the stack is running. Restart the SM node and stack.

can't set up my port, sleeping

- **SM Area**

  Discovery
- **Meaning**

  The SM cannot communicate with the stack.
- **Action**

  Make sure the stack is running. Restart the SM node and stack.

Get NodeInfo failed for local node. status 7

- **SM Area**

  Discovery
- **Meaning**

  The SM cannot communicate with the stack.
- **Action**

  If the condition persists, restart the SM node.

sm_setup_node: Get NodeDesc failed for local node, status 7

- **SM Area**

  Discovery
- **Meaning**

  The SM cannot communicate with the stack.
- **Action**

  If the condition persists, restart the SM node.

Error adding Node GUID: 0x00066a00d9000143 to tree. Already in tree!

- **SM Area**

  Discovery
- **Meaning**

  Duplicate Node GUID in fabric.
- **Action**

  Using fabric tools, locate the device with the duplicate node GUID and remove it.

Error adding Port GUID: 0x00066a00d9000143 to tree. Already in tree!

- **SM Area**

  Discovery
- **Meaning**

  Duplicate Port GUID in the fabric.
- **Action**

  Using fabric tools, locate the device with the duplicate port GUID and remove it.

Duplicate NodeGuid for Node HFI1 nodeType[1-3] guid 0x00066a00d9000143 and existing node[x] nodeType=1-3, HFI2, guid 0x00066a00d9000143

- **SM Area**

  Discovery
- **Meaning**

  A duplicate Node GUID in fabric.
- **Action**

  Using fabric tools, locate the device with the duplicate node GUID and remove it.

Marking port[x] of node[x] HFI1 guid 0x00066a00d9000143 DOWN in the topology

- **SM Area**

  Discovery

- **Meaning**

  Port x of HFI1 is not responding.

- **Action**

  Check the health of node HFI1.

Failed to init SL2SC/SC2SL Map (setting port down) on node HFI1/sw1 nodeGuid 0x00066a00d9000143 node index X port index Y

- **SM Area**

  Discovery

- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric. Persistence indicates that the node may be having problems.

- **Action**

  Check the health of node HFI1/SW1 if the fabric was idle or in persistent condition.

Failed to init VL Arb (setting port down) on node HFI1/Sw1 nodeGuid 0x00066a00d9000143 node index X port index Y

- **SM Area**

  Discovery

- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric. Persistence indicates that the node may be having problems.

- **Action**

  Check the health of node HFI1/SW1 if the fabric was idle or in persistent condition.

TT(ta): can't ARM/ACTIVATE node HFI1/sw1 guid 0x00066a00d9000143 node index X port index Y

- **SM Area**

  Discovery

- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric. Persistence indicates that the node may be having problems.

- **Action**

  Check the health of node HFI1/SW1 if the fabric was idle or in persistent condition.

sa_xxxxx: Reached size limit at X records

- **SM Area**

  Administrator

- **Meaning**

  The response buffer is too large.

- **Action**

  Contact support.

sa_NodeRecord_Set: NULL PORTGUID for Node Guid[0x00066a00d9000143], HFI1, Lid 0x1

- **SM Area**

  Administrator

- **Meaning**

  Possible data corruption.

- **Action**

  Contact support.

sa_TraceRecord: destination port is not in active state; port LID: 0x1 (port GUID 0x00066a00d9000144)

- **SM Area**

  Administrator

- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric.

- **Action**

  Check the health of the destination if the condition persists.

sa_TraceRecord: Cannot find path to port LID 0x2 (port guid 0x00066a00d9000144) from port LID 0x1 (port guid 0x00066a00d9000143)

- **SM Area**

  Administrator

- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric.

- **Action**

  Check for the next three messages.

sa_TraceRecord_Fill: Reached size limit while processing TRACE_RECORD request

- **SM Area**

  Administrator

- **Meaning**

  The response buffer is too large.

- **Action**

  Contact support.

sa_PathRecord: NULL PORTGUID in Source/Destination Gid 0xFE80000000000000:0000000000000000 of PATH request from Lid 0x1

- **SM Area**

  Administrator

- **Meaning**

  Invalid data in the SA request.

- **Action**

  Check the health of the requester at LID 0x1.

sa_PathRecord: Cannot find path to port LID 0x2 (port guid 0x00066a00d9000144) from port LID 0x1 (port guid 0x00066a00d9000143)

- **SM Area**

  Administrator

- **Meaning**

  May be caused by simultaneous removal/insertion events in the fabric.

- **Action**

  Check the health of the destination LID 0x2.

sa_PathRecord: Cannot find path to port LID 0x2 (port guid 0x00066a00d9000144) from port LID 0x1 (port guid 0x00066a00d9000143) with pkey 0x800d

- **SM Area**

  Administrator

- **Meaning**

  A path does not exist in the partition with the given PKey between the given source and destination.

- **Action**

  Check the configuration to determine if the path should exist in the given PKey. Check the health of the destination LID 0x2 if the configuration is valid.

sa_PathRecord: port LID 0x1 (port guid 0x00066a00d9000143) not a member of multicast group 0xff12401bffff0000:00000000ffffffff

- **SM Area**

  Administrator

- **Meaning**

  Group may have just been deleted or the requester is not a member of the group.

- **Action**

  None

sa_McMemberRecord_Set: Port GID in request (0x0080000000000000:0x0000000000000000)from HFI1, Port 0x00066a00d9000143, LID 0x1 has a NULL GUID/invalid prefix, returning status 0x0500

- **SM Area**

  Administrator

- **Meaning**

  Invalid data in the SA request.

- **Action**

  Check the health of HFI1.

sa_McMemberRecord_Set: MTU selector of 2 with MTU of 4 does not work with realizable MTU of 1 for request from compute-0-24, Port 0x00066A00A00005C5, LID 0x009C, returning status 0x0200

- **SM Area**

  Administrator

- **Meaning**

  The requester port data is not compatible with the group data.

- **Action**

  Create a group at the lowest common denominator or the host should join with a rate selector of “less than” rather than “exactly”.

sa_McMemberRecord_Set: Rate selector of 2 with Rate of 3 does not work with realizable Rate of 2 for request from compute-0-24, Port 0x00066A00A00005C5, LID 0x009C, returning status 0x0200

- **SM Area**

  Administrator

- **Meaning**

  Node compute-0-24 has requested a port rate that is incompatible with the group rate.

- **Action**

  Check that the requester port is not running at 1X width or that the multicast group was not created with a rate greater than what some of the host ports can support.

sa_McMemberRecord_Set: Component mask (0x00000000000XXXXX) does not have bits required to create (0x00000000000130C6) a group for new MGID of 0xFF12401BFFFF0000:00000000FFFFFFFF for request from HFI1, Port 0x00066a00d9000143

- **SM Area**

  Administrator

- **Meaning**

  End node may be trying to join a group that does not exist.

- **Action**

  OpenIB and Sun stacks require that the broadcast group be pre-created by the SM.

sa_McMemberRecord_Set: Component mask of 0x0000000000010083 does not have bits required (0x00000000000130C6) to CREATE a new group in request from HFI1, Port 0x00066a00d9000143

- **SM Area**

  Administrator

- **Meaning**

  Specific bits must be set in a CREATE group request.

- **Action**

  The requester is violating the protocol specification.

sa_McMemberRecord_Set: Bad (limited member) PKey of 0x1234 for request from ibhollab54 HFI-1, Port 0x00066a00d9000143, LID 0x1, returning status 0x200

- **SM Area**

  Administrator

- **Meaning**

  The PKey specified in request was limited, it should be full.

- **Action**

  Check configuration.

sa_McMemberRecord_Set: MC group create request denied for node ibhollab54 HFI-1, port 0x00066a00d9000144 from lid 0x2, failed VF validation (mgid=0xFF12401BFFFF0000:0x0000000000000016, sl=x, pkey=0xnnnn)

- **SM Area**

  Administrator

- **Meaning**

  An attempt to create a multicast group failed due to validation failures.

- **Action**

  Check configuration if create by source should be valid.

sa_McMemberRecord_Set: Invalid MGID (0xFF270000FFFF0000:00000000FFFFFFFF) in CREATE/JOIN request from HFI1, Port 0x00066a00d9000143, LID 0x0001, returning status 0x0500

- **SM Area**

  Administrator

- **Meaning**

  MGID requested violating the protocol specification.

- **Action**

  The requester is violating the protocol specification.

sa_McMemberRecord_Set: Join state of 0x1-2 not full member for NULL/NEW GID request from HFI1, Port 0x00066a00d9000143, LID 0x0001, returning status 0x0200

- **SM Area**

  Administrator

- **Meaning**

  Creation of a Multicast Group requires FULL membership.

- **Action**

  The requester is violating the protocol specification.

sa_McMemberRecord_Set: Join state of ~0x1 not full member for request to CREATE existing MGID of 0xFF12401BFFFF0000:00000000FFFFFFFF

- **SM Area**

  Administrator

- **Meaning**

  Creation of a Multicast Group requires FULL membership.

- **Action**

  The requester is violating the protocol specification.

sa_McMemberRecord_Set: Component mask of 0x00000000000XXXXX does not have bits required (0x0000000000010083) to JOIN group with MGID 0xFF12401BFFFF0000:00000000FFFFFFFF in request from %s, Port 0x%.16"CS64"X, LID 0x%.4X, returning status 0x%.4X

- **SM Area**

  Administrator

- **Meaning**

  Specific bits must be set in a JOIN group request.

- **Action**

  The requester is violating the protocol specification.

sa_McMemberRecord_Set: Maximum number groups reached (1000), failing CREATE request from HFI1, Port 0x00066a00d9000143, LID 0x0011, returning status 0x0100

- **SM Area**

  Administrator

- **Meaning**

  No resources.

- **Action**

  Delete some of the multicast groups or configure the SM to overload MLIDs during a group creation.

sa_McMemberRecord_Set: Failed to assign GID for CREATE request from HFI1, Port 0x00066a00d9000143, LID 0x0001, returning status 0x0100

- **SM Area**

  Administrator

- **Meaning**

  No resources.

- **Action**

  Delete some of the multicast groups or configure the SM to overload the MLIDs during group creation.

sa_McMemberRecord_Set: No multicast LIDs available for request from HFI1, Port 0x00066a00d9000143, LID 0x0001, returning status 0x0100

- **SM Area**

  Administrator

- **Meaning**

  No resources.

- **Action**

  Delete some of the multicast groups or configure the SM to overload the MLIDs during group creation.

sa_McMemberRecord_Set: MGID 0xFF12401BFFFF0000:00000000FFFFFFFF does not exist; Failing JOIN request from HFI1, Port 0x00066a00d9000143, LID 0x0001, returning status 0x0200

- **SM Area**

  Administrator

- **Meaning**

  JOIN of a group that does not exist.

- **Action**

  OpenIB and Sun stacks require that the broadcast group be pre-created by SM.

sa_McMemberRecord_Set: Qkey/PKey of 0x1234 does not match group QKey of 0x4321 for group 0xFF12401BFFFF0000:00000000FFFFFFFF for request from HFI1, Port 0x00066a00d9000143, LID 0x0001, returning status 0x0200

- **SM Area**

  Administrator

- **Meaning**

  SM may have been set to create the default broadcast group with parameters not valid for the fabric.

- **Action**

  Reconfigure the default broadcast group with the proper parameters.

sa_McMemberRecord_Set: Group MTU of 5 greater than requester port mtu of 2/4 for group 0xFF12401BFFFF0000:00000000FFFFFFFF for request from HFI1, Port 0x00066a00d9000143, LID 0x0001, returning status 0x0200

- **SM Area**

  Administrator

- **Meaning**

  SM may have been set to create the default broadcast group with parameters not valid for fabric.

- **Action**

  Reconfigure the default broadcast group with the proper parameters.

sa_McMemberRecord_Set: Group Rate/MTU of X is too low/high for requested rate/mtu of Y, rate/mtu selector of 2, and port rate/mtu of Z for group 0xFF12401BFFFF0000:00000000FFFFFFFF in request from HFI1

- **SM Area**

  Administrator

- **Meaning**

  The SM may have been set to create a default broadcast group with parameters not valid for a fabric or a host has created the group at a RATE not supported by other hosts.

- **Action**

  Create the group at the lowest common denominator or the host should join with a rate selector of “less than” rather than “exactly”.

sa_InformInfo: Subscription for security trap not from trusted source[lid=0x0001], smkey=0x0, returning status 0x0200

- **SM Area**

  Administrator

- **Meaning**

  Requester using wrong smkey.

- **Action**

  Make sure to use the same smkey as what is configured in the SM.

sm_resolve_pkeys_for_vfs: VFabric has application SA selected, bad PKey configured 0x1, must use Default PKey.",

- **SM Area**

  Administrator

- **Meaning**

  The SA select is limited to Virtual Fabrics using the Default PKey 0x7fff.

- **Action**

  Configuration change needed for SA Select.

sm_resolve_pkeys_for_vfs: Virtual Fabric VF0013 MulticastGroup configuration error, MGID does not match app, disabling Default Group

- **SM Area**

  Administrator

- **Meaning**

  The Multicast Group has an MGID configured that does not match any application that is part of this Virtual Fabric.

- **Action**

  Configuration change needed for multicast group creation.

sm_resolve_pkeys_for_vfs: Virtual Fabric VF0013 MulticastGroup configuration error, mismatch on pkey. Disabling Default Group

- **SM Area**

  Administrator

- **Meaning**

  The MulticastGroup linked to this Virtual Fabric do not share a common Pkey. Disabling the multicast group for this vFabric.

- **Action**

  Configuration change required if mcast group default creation is needed.

sm_initialize_port/sm_dbsync: cannot refresh sm PKeys

- **SM Area**

  Administrator

- **Meaning**

  An internal error occurred when attempting to refresh the SM PKeys.

- **Action**

  Contact customer support if condition persists.

sa_ServiceRecord_Add: Failed to ADD serviced record ID=0x1000000000003531 from lid 0x2 due to invalid pkey

- **SM Area**

  Administrator

- **Meaning**

  Serviced record add failure due to request with invalid PKey.

- **Action**

  Configuration change required if request should be granted.

mismatch from request port

- **SM Area**

  Administrator

- **Meaning**

  Serviced record add failure due to request with PKey not shared by requestor.

- **Action**

  Configuration change required if request should be granted.

sa_ServiceRecord_Add: Failed to ADD serviced record ID=0x1000000000003531 from lid 0x2 due to pkey mismatch from service port

- **SM Area**

  Administrator

- **Meaning**

  Serviced record add failure due to request with PKey not shared by requestor.

- **Action**

  Configuration change required if request should be granted.

4.3.4. IPoIB Troubleshooting

4.3.4.1. Viewing IPoIB Hardware Addresses

The `arp` command and `/proc/net/arp` do not show the full 20-byte IPoIB hardware address even when the IPoIB and hfi1 modules are working.

- `arp` shows only six bytes of the address.
- `/proc/net/arp` shows only nine bytes of the address.

To view the full 20-byte hardware address, use **`ip neigh show`**.

4.3.4.1.1. Address Details

The 20-byte IPoIB hardware MAC address is composed of the following (from most significant bit to least significant bit):

- 1 byte: Reserved (zero)
- 3 bytes: Unreliable datagram queue pair (UD QP) number assigned to IPoIB on a given endpoint
- 16 bytes: GID assigned to Port

  The 16-byte GID includes:

  – 8 bytes: Subnet prefix for CN5000 Omni-Path Fabric

  – 8 bytes: Port GUID

4.4. Performance Troubleshooting

Refer to the *CN5000 Performance Tuning Guide* for details about optimizing CN5000 Omni-Path Fabric performance and handling performance issues.

5. Working with Cornelis Technical Support

Cornelis Technical Support provides expert assistance for a variety of services including technical issues and questions, warranty, and returns.

Resources

- For product information, visit the Cornelis Networks Website.
- All downloadable documentation, software, and firmware packages are available through the Cornelis Customer Center.
- For online end-user publications, see the Cornelis Documentation Portal.
- For support information, visit the Cornelis Networks Support & Professional Services.
- For questions or requests, you can open a support case by sending an email to support@cornelisnetworks.com.

5.1. Technical Issues

If you are experiencing technical issues, the following methods are available to contact Cornelis Technical Support.

- Call Support (+1 (484) 497-9665): When you need to speak to an expert for a critical issue, we are available 24/7 by phone. Our Technical Support team is ready to address your issue at any time, day or night.
- Email Support: Cornelis Networks places customer satisfaction as a top priority. We are available 24 hours a day, 365 days a year.
- Web Request (Submit a Ticket): Our interactive, automated system tracks your support requests. You can submit a new request any time, day or night, and check on current projects, whether initiated by a phone call, email, or through this system.

Tips for Requesting Support

1. Open one case for each issue, using email or web request to allow a complete summary of the issue being reported.
2. Indicate urgency:
   - Describe your assessment of the severity/criticality.
   - Include the business/project impact.
   - Provide the current state of the system/issue: back in production, partially usable, or problem no longer exists.
   - If critical for your business and needing immediate attention, it is best to open a case through email or web request, then use the phone to alert the Cornelis Technical Support team of the critical issue.

3. Title/Subject should capture a concise summary of what problem is being seen:
   - Describe the main issue.
   - If Severity 1/Critical, add [Urgent] to the title.
4. Description should provide sufficient detail to start analysis:
   - What is the issue: failure, error messages, low performance, or other?
   - Were there any changes to the system leading up to this behavior?
   - What is the frequency of occurrence: ongoing, intermittent, triggered by, onetime, or other?
   - Include steps to reproduce the issue.
   - If support is in partnership with an OEM/vendor, has a case been opened with them? If so, provide the case number?
5. Provide logs and/or captures from the suspect hardware, links, and so on.

Appendix A. Glossary of Acronyms

The following acronyms are used throughout the Omni-Path Documentation Set.

ACPI	Advanced Configuration and Power Interface - This industry standard is an interface that allows the BIOS, OS, and peripherals to communicate with each other about power usage.
AOC	Active Optical Cable - This electrical-to-optical cable is used to connect modules in the CN5000 Omni-Path Fabric.
ASIC	Application-Specific Integrated Circuit - This is an integrated circuit chip used in Omni-Path Switches and SuperNICs.
DCS	Director Class Switch - This is a high-performance, scalable switch designed for interconnecting mid-sized and large High-Performance Computing (HPC) clusters, offering high bandwidth and low latency.
DLID	Destination Local Identifier - This is the address within the header of a data packet that identifies the destination computing node within the fabric.
FE	Fabric Executive - This is an Omni-Path component that provides out-of-band access to the Fabric Manager.
FM	Fabric Manager - This is an OPX Software component responsible for managing the fabric using management packets over a dedicated virtual lane (VL15).
FRU	Field Replaceable Unit - This is a part or assembly that can be quickly and easily replaced by the user or technician without having to send the product or system to a repair facility.
GID	Group Identifier - This is a unique identifier assigned to a group of users on a system, allowing them to access files and directories with common permissions.
GUID	Globally Unique Identifier - This is a 128-bit text string generated when a unique reference number is needed to identify a component (for example, hardware, software, node) on a computer or network.
IPoIB	Internet Protocol over InfiniBand- This network interface implementation is IP datagrams over the InfiniBand.

ISL	Inter-Switch Link - This is a Cisco Systems-proprietary protocol that maintains VLAN information in Ethernet frames as traffic flows between switches and routers, or switches and switches.
LID	Local Identifier - This is a 24-bit address local to a subnet that is assigned to end nodes by the subnet manager, providing additional features related to routing algorithms.
LQI	Link Quality Indicator - This indicator provides a value measuring the quality of received data packets. It can identify secure transmissions and links with poor signal integrity.
MGID	Multicast Group Identifier - This ID number maps multicast source, the multicast address, and the VLAN.
MPI	Message Passing Interface - This is a standardized means of exchanging messages between multiple computers running a parallel program across distributed memory
MTU	Maximum Transmission Unit - This is the size of the largest protocol data unit or data packet that can be sent in a single network layer transaction.
OUI	Organizationally Unique Identifier - Administered by the IEEE, this is the part of the MAC Address that identifies the vendor of the network adapter (often referred to as the *company_id*). It is the first three bytes of the six-byte field.
PA	Performance Administrator/Administration - This component allows a centralized control point for querying the performance data in the fabric.
PKey	Partition Key - This is a unique ID assigned to an InfiniBand or vFabric partition.
PM	Performance Manager/Performance Management - This is the Fabric Management entity responsible for monitoring fabric information related to the port and virtual lane level counters and the picture that they convey.
PMA	Performance Management Agent
QSFP	Quad Small Form-factor Pluggable - This is a compact, hot-pluggable transceiver used for high speed data communications applications.
SA	Subnet Administrator/Administration
SC	Service Channel

SI	Signal Integrity - This is a set of measures of the quality of an electrical signal.
SLID	Source Local Identifier
SM	Subnet Manager or Subnet Management
SMA	Subnet Management Agent
SNMP	Simple Network Management Protocol - This is an Internet Standard protocol for collecting and organizing information about managed devices on IP networks and for modifying that information to change device behavior.
VL	Virtual Lane - This feature allows multiple, multiplexed, independent data to stream onto a single physical link.